arXiv:1509.06390v1 [cs.DB] 21 Sep 2015

Exchange-Repairs: Managing Inconsistency in Data Exchange

Balder ten Cate
btencate@ucsc.edu
UC Santa Cruz
Google,

Richard L. Halpert
rhalpert@ucsc.edu
UC Santa Cruz

Phokion G. Kolaitis
kolaitis@ucsc.edu
UC Santa Cruz
IBM Research - Almaden

September 23, 2015

Abstract

In a data exchange setting with target constraints, it is often the case that a given source instance has no solutions. Intuitively, this happens when data sources contain inconsistent or conflicting information that is exposed by the target constraints at hand. In such cases, the semantics of target queries trivialize, because the certain answers of every target query over the given source instance evaluate to "true". The aim of this paper is to introduce and explore a new framework that gives meaningful semantics in such cases by using the notion of exchange-repairs. Informally, an exchange-repair of a source instance is another source instance that differs minimally from the first, but has a solution. In turn, exchange-repairs give rise to a natural notion of exchange-repair certain answers (in short, XR-certain answers) for target queries in the context of data exchange with target constraints.

After exploring the structural properties of exchange-repairs, we focus on the problem of computing the XR-certain answers of conjunctive queries. We show that for schema mappings specified by source-to-target GAV dependencies and target equality-generating dependencies (egds), the XR-certain answers of a target conjunctive query can be rewritten as the consistent answers (in the sense of standard database repairs) of a union of conjunctive queries over the source schema with respect to a set of egds over the source schema, thus making it possible to use a consistent query-answering system to compute XR-certain answers in data exchange. In contrast, we show that this type of rewriting is not possible for schema mappings specified by source-to-target LAV dependencies and target egds, nor for schema mappings specified by both source-to-target and target GAV dependencies. We then examine the general case of schema mappings specified by source-to-target GLAV constraints, a weakly acyclic set of target tgds and a set of target egds. The main result asserts that, for such settings, the XR-certain answers of conjunctive queries can be rewritten as the certain answers of a union of conjunctive queries with respect to the stable models of a disjunctive logic program over a suitable expansion of the source schema.

1 Introduction and Summary of Contributions

Data exchange is the problem of transforming data structured under one schema, called the source schema, into data structured under a different schema, called the target schema, in such a way that pre-specified constraints on these two schemas are satisfied. Data exchange is a ubiquitous data inter-operability task that has been explored in depth during the past decade (see [3]). This task is formalized with the aid of schema mappings $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{\mathrm{st}}, \Sigma_{\mathrm{t}})$, where $\mathbf{S}$ is the source schema, $\mathbf{T}$ is

the target schema, $\Sigma_{\rm st}$ is a set of constraints between $\mathbf{S}$ and $\mathbf{T}$, and $\Sigma_{\rm t}$ is a set of constraints on $\mathbf{T}$. The most thoroughly investigated schema mappings are the ones in which $\Sigma_{\rm st}$ is a set of source-to-target tuple-generating dependencies (s-t tgds) and $\Sigma_{\rm t}$ is a set of target tuple-generating dependencies (target tgds) and target equality-generating dependencies (target egds) [19]. An example of such a schema mapping, along with a target query, follows:

Every schema mapping $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{\rm st}, \Sigma_{\rm t})$ gives rise to two distinct algorithmic problems. The first is the existence and construction of solutions: given a source instance I, determine whether a *solution* for I exists (i.e., a target instance J so that (I, J) satisfies $\Sigma_{st} \cup \Sigma_t$) and, if it does, construct such a "good" solution. The second is to compute the *certain answers* of target queries, where if q is a target query and I is a source instance, then $\text{certain}(q, I, \mathcal{M})$ is the intersection of the sets $q(J)$, as J varies over all solutions for I. For arbitrary schema mappings specified by s-t tgds and target tgds and egds, both these problems can be undecidable [26]. However, as shown in [19], if the set $\Sigma_{\rm t}$ of target tgds obeys a mild structural condition, called *weak acyclicity*, then both these problems can be solved in polynomial time using the *chase procedure*. Given a source instance I, the chase procedure attempts to build a "most general" solution J for I by generating facts that satisfy each s-t tgd and each target tgd as needed, and by equating two nulls or equating a null to a constant, as dictated by the egds. If the chase procedure encounters an egd that equates two distinct constants, then it terminates and reports that no solution for I exists. Otherwise, it constructs a *universal* solution J for I, which can also be used to compute the certain answers of conjunctive queries in time bounded by a polynomial in the size of I.

Consider the situation in which the chase terminates and reports that no solution exists. In such cases, for every boolean target query q, the certain answers $\text{certain}(q, I, \mathcal{M})$ evaluate to "true". Even though the certain answers have become the standard semantics of queries in the data exchange context, there is clearly something unsatisfactory about this state of affairs, since the certain answers trivialize when no solutions exist. Intuitively, the root cause for the lack of solutions is that the source instance contains inconsistent or conflicting information that is exposed by the target constraints of the schema mapping at hand. In turn, this suggests that alternative semantics for target queries could be obtained by adopting the notions of database *repairs* and *consistent answers* from the study of inconsistent databases (see [7] for an overview). We note that several different types of repairs have been studied in the context of inconsistent databases; the most widely used ones are the *symmetric difference ($\oplus$-repairs)*, which contain as special cases the *subset-repairs* and the *superset-repairs*.

How can the notions of database repairs and consistent answers be adapted to the data exchange framework? When one reflects on this question, then one realizes that several different approaches are possible.

One approach, which we call *materialize-then-repair*, is as follows: given a source instance, a target instance is produced by chasing with the source-to-target tgds in $\Sigma_{\rm st}$ and the target tgds in $\Sigma_{\rm t}$, while ignoring the target egds in $\Sigma_{\rm t}$. Since the target instance produced this way may very well violate the egds in $\Sigma_{\rm t}$, it is treated as an inconsistent instance w.r.t. $\Sigma_{\rm t}$; consider its repairs. Note that a similar approach has been adopted by [8, 12] in the context of data integration. A different approach, which we call *exchange-as-repair*, treats the given source instance as an inconsistent instance over the combined schema $\mathbf{S} \cup \mathbf{T}$ w.r.t. the union $\Sigma_{\rm st} \cup \Sigma_{\rm t}$ and considers its repairs. Note that this is in the spirit of [24], where instances in peer data exchange that do not satisfy the schema mapping at hand are treated as inconsistent databases over a combined schema. We now point out that neither of these approaches gives rise to satisfactory semantics.

Figure 2 gives an example of a target instance that is produced in the materialize-then-repair approach by chasing with the s-t tgds in Figure 1. Clearly, J is inconsistent because it violates the egd in $\Sigma_{\rm t}$. Consider now the subset repair J' in Figure 3 of our materialized target instance J (note that, in this case, symmetric difference repairs coincide with subset repairs). Notice that the repair J' places peter in the exec department, yet still has him performing tasks for the software department – the fact that the

$$\Sigma_{\text{st}} = \left\{ \begin{array}{lcl} \texttt{Task_Assignments}(p,t,d) & \rightarrow & \texttt{Departments}(p,d) \wedge \texttt{Tasks}(p,t) \\ \texttt{Stakeholders_old}(t,s) & \rightarrow & \texttt{Stakeholders_new}(t,s) \end{array} \right\}$$

$$\Sigma_{\text{t}} = \left\{ \texttt{Departments}(p,d) \wedge \texttt{Departments}(p,d') \rightarrow d = d' \right\}$$

$$\texttt{boss}(person, stakeholder) = \exists task. \\ \texttt{Tasks}(person, task) \wedge \texttt{Stakeholders_new}(task, stakeholder)$$

Figure 1: A schema mapping $\mathcal{M}$ specified by tgds and egds, and a target query. In this example, the egd is actually a key constraint and there are no target tgds.

I

Task_Assignments		
person	task	dept
peter	tpsreport	software
peter	spaceout	software
peter	meetbobs	exec

Stakeholders_old	
task	stakeholder
tpsreport	lumbergh
tpsreport	portman
spaceout	bobs
meetbobs	bobs

J

Departments	
person	dept
peter	software
peter	exec

Tasks	
person	task
peter	tpsreport
peter	spaceout
peter	meetbobs

Stakeholders_new	
task	stakeholder
tpsreport	lumbergh
tpsreport	portman
spaceout	bobs
meetbobs	bobs

Figure 2: A source instance I and the inconsistent target instance J that results from chasing I with the tgds in $\mathcal{M}$.

J′

Departments	
person	dept
~~peter~~	~~software~~
peter	exec

Tasks	
person	task
peter	tpsreport
peter	spaceout
peter	meetbobs

Stakeholders_new	
task	stakeholder
tpsreport	lumbergh
tpsreport	portman
spaceout	bobs
meetbobs	bobs

Figure 3: A symmetric-difference-repair of J w.r.t. Σ_{t}.

"tpsreport" and "spaceout" tasks are derived from a tuple placing peter in the software department has been lost. The only other repair of J similarly fails to reflect the shared origin of tuples in the `Tasks` and `Departments` tables, and this disconnect in the materialize-then-repair approach manifests in the consistent answers to target queries. In this example, the consistent answers for $\texttt{boss}(\text{peter}, b)$ are $\{(\text{peter}, \text{bobs}), (\text{peter}, \text{portman}), (\text{peter}, \text{lumbergh})\}$. However, the last two tuples are derived from facts placing peter in the software department, even though in J' he is not.

The situation is no better in the exchange-as-repair approach. Figure 4 depicts three repairs of this type (using symmetric difference semantics). While the first two repairs in Figure 4 seem reasonable, in the third we have eliminated `Task_Assignments`(peter, spaceout, software), even though our key constraint is already satisfied by the removal of `Task_Assignments`(peter, meetbobs, exec) alone. In this approach, the consistent answers of $\texttt{boss}(\text{peter}, b)$ are $\emptyset$, despite the intuitive conclusion that peter should be performing tasks for the bobs regardless of which way we fix the department key constraint violation. For symmetric-difference repairs, it is equally valid to satisfy a violated tgd by removing tuples as by adding them[1]. However, in a data exchange setting, the target instance is initially empty, so it would be more natural to satisfy violated tgds by deriving new tuples. This observation motivates the particulars of our approach, which we

[1] A noteworthy alternative to symmetric difference repairs are the loosely-sound semantics of [14], discussed in detail in Section 3.3.

(I_1, J_1)

Task_Assignments		
person	task	dept
~~peter~~	~~tpsreport~~	~~software~~
~~peter~~	~~spaceout~~	~~software~~
peter	meetbobs	exec

Departments	
person	dept
peter	exec

Tasks	
person	task
peter	meetbobs

(I_2, J_2)

Task_Assignments		
person	task	dept
peter	tpsreport	software
peter	spaceout	software
~~peter~~	~~meetbobs~~	~~exec~~

Departments	
person	dept
peter	software

Tasks	
person	task
peter	tpsreport
peter	spaceout

(I_3, J_3)

Task_Assignments		
person	task	dept
peter	tpsreport	software
~~peter~~	~~spaceout~~	~~software~~
~~peter~~	~~meetbobs~~	~~exec~~

Departments	
person	dept
peter	software

Tasks	
person	task
peter	tpsreport

Figure 4: Three repairs of $(I, \emptyset)$ (from Figure 2) w.r.t. $\Sigma_{\mathrm{st}} \cup \Sigma_{\mathrm{t}}$. The `Stakeholders` tables are omitted for brevity.

introduce next.

1.1 Summary of Contributions

Our aim in this paper is to introduce and explore a new framework that gives meaningful and non-trivial semantics to queries in data exchange, including cases in which no solutions exist for a given source instance.

At the conceptual level, the main contribution is the introduction of the notion of an *exchange-repair*. Informally, an exchange-repair of a source instance is another source instance that differs minimally from the first, but has a solution. Exchange-repairs give rise to a natural notion of *exchange-repair certain answers* (in short, *XR-certain answers*) for target queries in the context of data exchange. Note that if a source instance I has a solution, then the XR-certain answers of target queries on I coincide with the certain answers of the queries on I. If I has no solutions, then unlike the certain answers, the XR-certain answers are non-trivial and meaningful.

We provide examples demonstrating that these new semantics improve upon both the materialize-then-repair approach and the exchange-as-repair approach discussed earlier. We also produce a detailed comparison of the XR-certain semantics with the main notions of inconsistency-tolerant semantics studied in data integration and in ontology-based data access. This comparison is carried out in Section 3.2, after we have introduced our framework and presented some basic structural properties of exchange-repairs in Section 3.

After this, we focus on the problem of computing the XR-certain answers of conjunctive queries. In Section 4, we show that for schema mappings specified by source-to-target GAV (global-as-view) dependencies and target egds, the XR-certain answers of conjunctive queries can be rewritten as the consistent answers (in the sense of standard database repairs) of a union of conjunctive queries over the source schema with respect to a set of egds over the source schema, thus making it possible to use a consistent query-answering system to compute XR-certain answers in data exchange. In contrast, we show that this type of rewriting is not possible for schema mappings specified by source-to-target LAV (local-as-view) dependencies and target egds, nor for schema mappings specified by source-to-target and target GAV dependencies and target egds.

In Section 5, we examine the general case of schema mappings specified by s-t tgds, a weakly acyclic set of target tgds and a set of target egds. The main result is that, for such settings, the XR-certain answers of conjunctive queries can be rewritten as the certain answers of a union of conjunctive queries with respect to the stable models of a disjunctive logic program over a suitable expansion of the source schema. This is achieved in two steps. First, for schema mappings consisting of GAV s-t tgds, GAV target tgds, and target egds, we show that the XR-certain answers of conjunctive queries can be reduced to cautious reasoning over stable models of a disjunctive

logic program. Second, for schema mappings consisting of GLAV s-t tgds, weakly acyclic sets of GLAV target tgds, and target egds, we show that the XR-certain answers of conjunctive queries can be rewritten as the XR-certain answers of conjunctive queries w.r.t. a schema mapping consisting of GAV s-t tgds, GAV target tgds, and target egds. In fact, we prove the stronger result that such a rewriting is possible for schema mappings specified by a second-order s-t tgd, a weakly acyclic second-order target tgd, and set of target egds.

2 Preliminaries

This section contains definitions of basic notions and a minimum amount of background material. Detailed information about schema mappings and certain answers can be found in [3, 19], and about repairs and consistent answers in [4, 7].

2.1 Instances and Homomorphisms

Fix an infinite set Const of elements, and an infinite set Nulls of elements such that Const and Nulls are disjoint. A *schema* $\mathbf{R}$ is a finite set of relation symbols, each having a designated arity. An $\mathbf{R}$-*instance* is a finite database I over the schema $\mathbf{R}$ whose active domain is a subset of $\text{Const} \cup \text{Nulls}$. A *fact* of an $\mathbf{R}$-instance I is an expression of the form $R(a_1, \ldots, a_k)$, where R is a relation symbol of arity k in $\mathbf{R}$ and $(a_1, \ldots, a_k)$ is a member of the relation R^I on I that interprets the relation symbol R. Every $\mathbf{R}$-instance can be identified with the set of its facts. We say that an $\mathbf{R}$-instance I' is a *sub-instance* of an $\mathbf{R}$-instance I if $I' \subseteq I$, where I' and I are viewed as sets of facts.

By a *homomorphism* between two instances K and K', we mean a map from the active domain of K to the active domain of K' that is the identity function on all elements of Const and such that for every atom $\mathrm{R}(v_1, ..., v_n) \in K$ we have that $\mathrm{R}(h(v_1), ..., h(v_n)) \in K'$.

2.2 Schema Mappings and Certain Answers.

A *tuple-generating dependency (tgd)* is an expression of the form $\forall \mathbf{x}(\phi(\mathbf{x}) \rightarrow \exists \mathbf{y} \psi(\mathbf{x}, \mathbf{y}))$, where $\phi(\mathbf{x})$ and $\psi(\mathbf{x}, \mathbf{y})$ are conjunctions of atoms over some relational schema.

Tgds are also known as GLAV (global-and-local-as-view) constraints. Tgds with no existentially quantified variables are called *full*. Two important special cases are the GAV constraints and the LAV constraints: the former are the tgds of the form $\forall \mathbf{x}(\phi(\mathbf{x}) \rightarrow P(\mathbf{x}))$ and the latter are the tgds of the form $\forall \mathbf{x}(R(\mathbf{x}) \rightarrow \exists \mathbf{y} \psi(\mathbf{x}, \mathbf{y}))$, where P and R are individual relation symbols. Every full tgd is logically equivalent to a set of GAV tgds that can be computed in linear time.

Suppose we have two disjoint relational schemas $\mathbf{S}$ and $\mathbf{T}$, called the *source* schema and the *target* schema. A *source-to-target tgd* (*s-t tgd*) is a tgd as above such that $\phi(\mathbf{x})$ is a conjunction over $\mathbf{S}$ and $\psi(\mathbf{x}, \mathbf{y})$ is a conjunction over $\mathbf{T}$. When the schemas are understood from context, we may say just tgd even if the constraint is source-to-target.

An *equality-generating dependency (egd)* is an expression of the form $\forall \mathbf{x}(\phi(\mathbf{x}) \rightarrow x_i = x_j)$ with $\phi(\mathbf{x})$ a conjunction of atoms over a relational schema.

For the sake of readability, we will frequently drop universal quantifiers when writing tgds and egds.

A *schema mapping* is a quadruple $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{\text{st}}, \Sigma_{\text{t}})$, where $\mathbf{S}$ is a source schema, $\mathbf{T}$ is a target schema, Σ_{st} is a finite set of source-to-target constraints, and Σ_{t} is a finite set of constraints over the target schema.

We will use the notation GLAV, GAV, LAV, EGD to denote the classes of sets of constraints consisting of finite sets of, respectively, GLAV constraints, GAV constraints, LAV constraints, and egds. If C is a class of sets of source-to-target dependencies and D is a class of sets of target dependencies, then the notation $C{+}D$ denotes the class of all schema mappings $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{\text{st}}, \Sigma_{\text{t}})$ such that Σ_{st} is a member of C and Σ_{t} is a member of D. For example, GLAV+EGD denotes the class of all schema mappings $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{\text{st}}, \Sigma_{\text{t}})$ such that Σ_{st} is a finite set of s-t tgds and Σ_{t} is a finite set of egds. Moreover, we will

use the notation (D_1, D_2) to denote that the union of two classes D_1 and D_2 of sets of target dependencies. For example, GAV+(GAV, EGD) denotes the class of all schema mappings $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{\text{st}}, \Sigma_{\text{t}})$ such that Σ_{st} is a set of GAV s-t tgds and Σ_{t} is the union of a finite set of GAV target tgds with a finite set of target egds.

Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{\text{st}}, \Sigma_{\text{t}})$ be a schema mapping. A target instance J is a *solution* for a source instance I w.r.t. $\mathcal{M}$ if J is finite, and the pair (I, J) satisfies $\mathcal{M}$, i.e., I and J together satisfy Σ_{st}, and J satisfies Σ_{t}. Recall that, by definition, instances are finite. Additionally, by convention, we will assume that source instances do not contain null values. A *universal* solution for I is a solution J for I such that if J' is a solution for I, then there is a homomorphism h from J to J' that is the identity on the active domain of I. If $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{\text{st}}, \Sigma_{\text{t}})$ is an arbitrary schema mapping, then a given source instance may have no solution or it may have a solution, but no (finite) universal solution. However, if Σ_{t} is the union of a *weakly acyclic* set of target tgds and a set of egds, then a solution exists if and only if a universal solution exists. Moreover, the *chase procedure* can be used to determine if, given a source instance I, a solution for I exists and, if it does, to actually construct a universal solution $chase(I, \mathcal{M})$ for I in time polynomial in the size of I (see [19] for details). The definition of weak acyclicity is given next, followed by the definition of the *chase procedure*.

Definition 2.1 ([19])**.** Let Σ be a set of tgds over a schema $\mathbf{T}$. Construct a directed graph, called the *dependency graph*, as follows:

- Nodes: For every pair (R, A) with R a relation symbol in $\mathbf{T}$ and A an attribute of R, there is a distinct node; call such a pair (R, A) a *position*.
- Edges: For every tgd $\forall \mathbf{x}(\phi(\mathbf{x}) \rightarrow \exists \mathbf{y}\, \psi(\mathbf{x}, \mathbf{y}))$ in Σ and for every x in $\mathbf{x}$ that occurs in ψ, and for every occurrence of x in ϕ in position (R, A_i):
  1. For every occurrence of x in ψ in position (S, B_j), add an edge $(R, A_i) \rightarrow (S, B_j)$ (if it does not already exist).
  2. For every existentially quantified variable y and for every occurrence of y in ψ in position (T, C_k), add a *special edge* $(R, A_i) \rightarrow (T, C_k)$ (if it does not already exist).

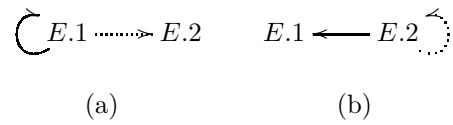

Figure 5: The dependency graphs for (a) $\forall x \forall y(E(x, y) \rightarrow \exists z\, E(x, z))$ and (b) $\forall x \forall y(E(x, y) \rightarrow \exists z\, E(y, z))$. Special edges are dotted.

We say that Σ is *weakly acyclic* if the dependency graph has no cycle going through a special edge.

WAGLAV denotes the class of all finite weakly acyclic sets of target tgds.

The tgd $\forall x \forall y(E(x, y) \rightarrow \exists z\, E(x, z))$ is weakly acyclic; in contrast, the tgd $\forall x \forall y(E(x, y) \rightarrow \exists z\, E(y, z))$ is not, because the dependency graph contains a special self-loop (see Figure 5). Moreover, every set of GAV tgds is weakly acyclic, since the dependency graph contains no special edges in this case.

What follows is the definition of the *chase procedure*.

Definition 2.2 (chase procedure [19])**.** Let K be an instance.

(tgd) Let d be a tgd $\phi(\mathbf{x}) \rightarrow \exists \mathbf{y} \psi(\mathbf{x}, \mathbf{y})$. Let h be a homomorphism from $\phi(\mathbf{x})$ to K such that there is no extension of h to a homomorphism h' from $\phi(\mathbf{x}) \wedge \psi(\mathbf{x}, \mathbf{y})$ to K. We say that d *can be applied to K with homomorphism h*.
Let K' be the union of K with the set of facts obtained by: (a) extending h to h' such that each variable in $\mathbf{y}$ is assigned a fresh labeled null, followed by (b) taking the image of the atoms of ψ under h'. We say that *the result of applying d to K with h* is K', and write $K \xrightarrow{d,h} K'$.

(egd) Let d be an egd $\phi(\mathbf{x}) \rightarrow (x_1 = x_2)$. Let h be a homomorphism from $\phi(\mathbf{x})$ to K such that $h(x_1) \neq h(x_2)$. We say that d *can be applied to K with homomorphism h*. We distinguish two cases.

- If both $h(x_1)$ and $h(x_2)$ are in $Const$ then we say that *the result of applying d to K with h* is "failure", and write $K \xrightarrow{d,h} \perp$.

- Otherwise, let K' be K where we identify $h(x_1)$ and $h(x_2)$ as follows: if one is a constant, then the labeled null is replaced everywhere by the constant; if both are labeled nulls, then one is replaced everywhere by the other. We say that *the result of applying d to K with h is K'*, and write $K \xrightarrow{d,h} K'$.

In the above, $K \xrightarrow{d,h} K'$ (including the case where K' is $\perp$) is called a *chase step*. We now define chase sequences and finite chases.

Let Σ be a set of tgds and egds, and let K be an instance.

- A *chase sequence of K with Σ* is a sequence (finite or infinite) of chase steps $K_i \xrightarrow{d_i,h_i} K_{i+1}$, with $i = 0, 1, ...$, with $K = K_0$ and d_i a dependency in Σ.
- A *finite chase of K with Σ* is a finite chase sequence $K_i \xrightarrow{d,h} K_{i+1}$, $0 \leq i < m$, with the requirement that either (a) $K_m = \perp$ or (b) there is no dependency d_i of Σ and there is no homomorphism h_i such that d_i can be applied to K_m with h_i. We say that K_m is the result of the finite chase. We refer to case (a) as the case of a *failing finite chase* and we refer to case (b) as the case of a *successful finite chase*.

In the context of data exchange, we chase the source instance first with the source-to-target constraints, and then continue chasing with the target constraints. The nature of s-t tgds ensure that no atoms are created over the source schema, so in this setting the result of chasing a source instance I with a schema mapping $\mathcal{M}$ is a pair (I, J) where J is a target instance. We usually refer to J alone as the result of the chase.

We will also make use of the notion of *rank* [19]. Let Σ be a finite weakly acyclic set of tgds. For every node (R, A) in the dependency graph of Σ, define an *incoming path* to be any (finite or infinite) path ending in (R, A). Define the *rank* of (R, A), denoted by $rank(R, A)$, as the maximum number of special edges on any such incoming path. Since Σ is weakly acyclic, there are no cycles going through special edges; hence, $rank(R, A)$ is finite. The *rank* of Σ, denoted $rank(\Sigma)$ is the maximum of $rank(R, A)$ over all positions (R, A) in the dependency graph of Σ.

If q is a query over the target schema $\mathbf{T}$ and I is a source instance, then the *certain answers* of q with respect to $\mathcal{M}$ are defined as

$$\text{certain}(q, I, \mathcal{M}) = \bigcap\{q(J) : \ J \text{ is a solution for } I \text{ w.r.t. } \mathcal{M}\}$$

Definition 2.3. Let J be an instance which may contain null values, and let q be a conjunctive query over the schema of J. Then $q{\downarrow}\,(J)$ is defined as the answers of q on J that contain no null values.

If J is a universal solution for a source instance I w.r.t. a schema mapping $\mathcal{M}$, then for every conjunctive query q, it holds that $\text{certain}(q, I, \mathcal{M}) = q{\downarrow}\,(J)$.

2.3 Repairs and Consistent Answers.

Let Σ be a set of constraints over some relational schema. An *inconsistent* database is a database that violates at least one constraint in Σ. Informally, a *repair* of an inconsistent database I is a consistent database I' that differs from I in a "minimal" way. This notion can be formalized in several different ways [4].

1. A *symmetric-difference-repair* of I, denoted $\oplus$-repair of I, is an instance I' that satisfies Σ and where there is no instance I'' such that $I \oplus I'' \subset I \oplus I'$ and I'' satisfies Σ. Here, $I \oplus I'$ denotes the set of facts that form the symmetric difference of the instances I and I'.

2. A *subset-repair* of I is an instance I' that satisfies Σ and where there is no instance I'' such that $I' \subset I'' \subseteq I$ and I'' satisfies Σ.

3. A *superset-repair* of I is an instance I' that satisfies Σ and where there is no instance I'' such that $I' \supset I'' \supseteq I$ and I'' satisfies Σ.

Clearly, subset-repair and superset-repairs are also $\oplus$-repairs; however, a $\oplus$-repair need not be a subset-repair or a superset-repair.

The *consistent answers* of a query q on I with respect to Σ are defined as:

$$\oplus\text{-CQA}(q, I, \Sigma) = \bigcap\{q(I') : \ I' \text{ is a } \oplus\text{-repair of } I \text{ w.r.t. } \Sigma\}$$

with subset and superset versions defined analogously.

3 Framework and Related Work

In this section, we introduce the exchange-repair framework, discuss its structural and algorithmic properties, and explore its relationship to inconsistency tolerant semantics in data integration and ontology-based data access.

3.1 The Exchange-Repair Framework

Definition 3.1. Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{\text{st}}, \Sigma_{\text{t}})$ be a schema mapping, I a source instance, and (I', J') a pair of a source instance and a target instance.

1. We say that (I', J') is a *symmetric-difference exchange-repair solution* (in short, a $\oplus$-XR-solution) for I w.r.t. $\mathcal{M}$ if (I', J') satisfies $\mathcal{M}$, and there is no pair of instances (I'', J'') such that $I \oplus I'' \subset I \oplus I'$ and (I'', J'') satisfies $\mathcal{M}$.
2. We say that (I', J') is a *subset exchange-repair solution* (in short, a subset-XR-solution) for I with respect to $\mathcal{M}$ if $I' \subseteq I$ and (I', J') satisfies $\mathcal{M}$; and there is no pair of instances (I'', J'') such that $I' \subset I'' \subseteq I$ and (I'', J'') satisfies $\mathcal{M}$.

Note that the minimality condition in the preceding definitions applies to the source instance I', but not to the target instance J' of the pair (I', J'). The source instance I' of a $\oplus$-XR-solution (subset-XR-solution) for I is called a $\oplus$-*source-repair* (respectively, subset source-repair) of I.

Figure 6 shows all two XR-solutions for our source instance and schema mapping. Notice that the shared origins of tuples are taken into account (for example, peter performs tasks only for his assigned department, unlike in Figure 3), but the XR-solutions retain more derived target information than the instances in Figure 4 (by preferring to satisfy tgds by adding rather than deleting tuples). If we now evaluate $\texttt{boss}(\text{peter}, b)$ over each target instance, and take the intersection, we have $\{(\text{peter}, \text{bobs})\}$, which aligns well with our intuitive expectations. A precise semantics for query answering is given later in this section.

Source-repairs constitute a new notion that, in general, has different properties from those of the standard database repairs. Indeed, as mentioned earlier, a $\oplus$-repair need not be a subset repair. In contrast, Theorem 3.2 (below) asserts that the state of affairs is different for source-repairs. Recall that, according to the notation introduced earlier, GLAV+(WAGLAV, EGD) denotes the collection of all schema mappings $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{\text{st}}, \Sigma_{\text{t}})$ such that Σ_{st} is a finite set of s-t tgds and Σ_{t} is the union of a finite weakly acyclic set of target tgds with a finite set of target egds.

Lemma 3.1. *Let $\mathcal{M}$ be a* GLAV+(GLAV, EGD) *schema mapping. If $I' \supseteq I$ are two source instances, then every solution for I' w.r.t. $\mathcal{M}$ is also a solution for I w.r.t. $\mathcal{M}$, and consequently if I has no solution w.r.t. $\mathcal{M}$ then I' has no solution w.r.t. $\mathcal{M}$.*

Proof. Let $I' \supseteq I$ be two source instances. We will show that if I' has a solution w.r.t. $\mathcal{M}$ then I also has a solution w.r.t. $\mathcal{M}$. Let J be an arbitrary solution for I' w.r.t. $\mathcal{M}$. Let $\phi(\mathbf{x}) \rightarrow \exists \mathbf{y} \psi(\mathbf{x}, \mathbf{y})$ be an arbitrary tgd in Σ_{st}, and let $h : \mathbf{x} \rightarrow \text{adom}(I)$ be a homomorphism such that $h(\phi(\mathbf{x})) \subseteq I$, and of course $h(\phi(\mathbf{x})) \subseteq I'$ as well. Then h can be extended to some homomorphism h' such that $h'(\psi(\mathbf{x}, \mathbf{y})) \subseteq J$, and therefore (I, J) together satisfy Σ_{st}, and since J satisfies Σ_{t}, we have that J is also a solution for I w.r.t. $\mathcal{M}$. □

Theorem 3.2. *Let $\mathcal{M}$ be a* GLAV+(GLAV, EGD) *schema mapping. Let I be a source instance. Then if (I', J') is a $\oplus$-XR-solution of I w.r.t. $\mathcal{M}$, then (I', J') is actually a subset-XR-solution of I w.r.t. $\mathcal{M}$. Consequently, every $\oplus$-source-repair of I is also a subset-source-repair of I.*

$(I, J)'$

Task_Assignments		
person	task	dept
~~peter~~	~~tpsreport~~	~~software~~
~~peter~~	~~spaceout~~	~~software~~
peter	meetbobs	exec

Departments	
person	dept
peter	exec

Tasks	
person	task
peter	meetbobs

$(I, J)''$

Task_Assignments		
person	task	dept
peter	tpsreport	software
peter	spaceout	software
~~peter~~	~~meetbobs~~	~~exec~~

Departments	
person	dept
peter	software

Tasks	
person	task
peter	tpsreport
peter	spaceout

Figure 6: Two XR-solutions for I w.r.t. $\mathcal{M}$. The `Stakeholders` tables are omitted for brevity.

Proof. Let (I', J') be an $\oplus$-XR-solution for I w.r.t. $\mathcal{M}$. Suppose $I' \setminus I \neq \emptyset$. Then by Lemma 3.1, J' is also a solution for $I' \cap I$. Since $I' \oplus I \supset (I' \cap I) \oplus I$, we have that (I', J') fails the minimality criterion and thus is not a $\oplus$-XR-solution for I w.r.t. $\mathcal{M}$, which is a contradiction. □

Remark. From here on and in view of Theorem 3.2, we will use the term XR-solution to mean subset-XR-solution; similarly, source-repair will mean subset source-repair.

Note that if $\mathcal{M}$ is a GLAV+(WAGLAV, EGD) schema mapping, then source-repairs always exist. The reason is that, since the pair $(\emptyset, \emptyset)$ trivially satisfies $\mathcal{M}$, then for every source instance I, there must exist a maximal subinstance I' of I for which a solution J' w.r.t. $\mathcal{M}$ exists; hence, (I', J') is a source repair for I w.r.t. $\mathcal{M}$.

We now claim that the following statements are true for arbitrary source instances and schema mappings.

1. Repairs of the target instance obtained by chasing with the tgds of the schema mapping are not necessarily XR-solutions.

2. Repairs of $(I, \emptyset)$ are not necessarily XR-solutions.

For the first statement, consider the pair (I, J') from Figures 2 and 3, where J' is a $\oplus$-repair of the inconsistent result J of the chase of I. Clearly, (I, J') is not an XR-solution, because J' is not a solution for I. For the second statement, consider the pairs (I_1, J_1), (I_2, J_2), (I_3, J_3) in Figure 4, all of which are $\oplus$-repairs of $(I, \emptyset)$. The first two are also XR-solutions of I, but the third one is not.

It can also be shown that XR-solutions are not necessarily $\oplus$-repairs of $(I, \emptyset)$. We now describe an important case in which XR-solutions **are** $\oplus$-repairs of $(I, \emptyset)$. For this, we recall the notion of a *core universal solution* from [19]. By definition, a core universal solution is a universal solution that has no homomorphism to a proper subinstance. If a universal solution exists, then a core universal solution also exists. Moreover, core universal solutions are unique up to isomorphism.

Proposition 3.3. *Let $\mathcal{M}$ be a* GLAV+(GLAV, EGD) *schema mapping. If I is source instance and (I', J') is an XR-solution for I w.r.t. $\mathcal{M}$ such that J' is a core universal solution for I' w.r.t. $\mathcal{M}$, then (I', J') is a $\oplus$-repair of $(I, \emptyset)$ w.r.t. $\Sigma_{st} \cup \Sigma_t$.*

Proof. Let (I'', J'') be a pair of instances which together satisfy $\Sigma_{\mathrm{st}} \cup \Sigma_{\mathrm{t}}$ and such that $(I'', J'') \oplus (I, \emptyset) \subseteq (I', J') \oplus (I, \emptyset)$. Since (I', J') is an XR-solution for I w.r.t. $\mathcal{M}$, there is no instance I'' such that $I' \subset I'' \subseteq I$ and I'' has a solution w.r.t. $\mathcal{M}$. Therefore it must be that $I'' = I'$. Furthermore, since J' is a core universal solution, there is no proper subinstance $J'' \subset J'$ that is a solution for I' w.r.t. $\mathcal{M}$, so $J'' = J'$. Therefore (I', J') is a $\oplus$-repair of $(I, \emptyset)$ w.r.t. $\Sigma_{\mathrm{st}} \cup \Sigma_{\mathrm{t}}$. □

Next, we present the second key notion in the exchange-repair framework.

Definition 3.2. Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{\text{st}}, \Sigma_{\text{t}})$ be a schema mapping and q a query over the target schema $\mathbf{T}$. If I is a source instance, then the *XR-certain answers* of q on I w.r.t. $\mathcal{M}$ is the set

$$\text{XR-certain}(q, I, \mathcal{M}) = \bigcap\{q(J') : (I', J') \text{ is an XR-solution for } I\}.$$

Note that when I has a solution w.r.t. M, it is its own only XR-solution. Thus the XR-certain semantics coincide with certain semantics when solutions exist. The next results provide a comparison of the XR-certain answers with the consistent answers.

Proposition 3.4. *Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{st}, \Sigma_t)$ be a* GLAV+(WAGLAV, EGD) *schema mapping and q a conjunctive query over the target schema $\mathbf{T}$. If I is a source instance, then* XR-certain$(q, I, \mathcal{M}) \supseteq \oplus$-CQA$(q, (I, \emptyset), \Sigma_{st} \cup \Sigma_t)$. *Moreover, this containment may be a proper one.*

Proof. Since $\mathcal{M}$ is weakly acyclic, for any instance I for which solutions exist, a core universal solution also exists. Therefore, we have that XR-certain$(q, I, \mathcal{M}) = \bigcap\{q(J') : (I', J')$ is an XR-solution for I w.r.t. $\mathcal{M}$, and J' is a core universal solution for I' w.r.t. $\mathcal{M}$ $\}$. By Proposition 3.3, the set of XR-solutions (I', J') where J' is a core universal solution for I' w.r.t. $\mathcal{M}$ is a subset (maybe proper) of the set of $\oplus$-repairs of $(I, \emptyset)$ w.r.t. $\Sigma_{\text{st}} \cup \Sigma_{\text{t}}$. Therefore XR-certain$(q, I, \mathcal{M}) \supseteq \oplus$-CQA$(q, (I, \emptyset), \Sigma_{\text{st}} \cup \Sigma_{\text{t}})$.

To see that this containment may be a proper one, consider the schema mapping $\mathcal{M}$ and query $\texttt{boss}(\texttt{peter}, b)$ in Figure 1, and the repairs of $(I, \emptyset)$ in Figure 4. It is easy to verify that $\oplus$-CQA$(\texttt{boss}(\texttt{peter}, b), (I, \emptyset), \Sigma_{\text{st}} \cup \Sigma_{\text{t}}) = \emptyset$, while XR-certain$(\texttt{boss}(\texttt{peter}, b), I, \mathcal{M}) = \{(\texttt{peter}, \texttt{bobs})\}$. □

The following proposition pertains to the case where Σ_{st} is the *copy mapping*, i.e. for each relation $R \in \mathbf{S}$ there is a corresponding relation R' of the same arity in $\mathbf{T}$, and Σ_{st} contains only the tgd $R(\mathbf{x}) \rightarrow R'(\mathbf{x})$ for each $R \in \mathbf{S}$. We say an instance J is the *copy* of an instance I if J is the canonical universal solution for I w.r.t. the copy mapping (so it contains the same facts up to renaming of relations).

Proposition 3.5. *Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{st}, \Sigma_t)$ be a* GAV+EGD *schema mapping where Σ_{st} is the copy mapping, and let q be a conjunctive query over the target schema $\mathbf{T}$. Then for every instance I, it holds that* XR-certain$(q, I, \mathcal{M})$ = *subset*-CQA(q, J, Σ_t), *where J is the copy of I.*

Proof. Since Σ_{st} specifies the copy mapping and Σ_{t} contains only egds, for every source repair I' there is an XR-solution(I', J') where J' is the copy of I'. Furthermore, J' is a universal solution for I' w.r.t. $\mathcal{M}$, so we can write XR-certain$(q, I, \mathcal{M}) = \bigcap\{q(J') \mid (I', J')$ is an XR-solution for I w.r.t. $\mathcal{M}$ and J' is the copy of $I'\}$. Therefore XR-certain$(q, I, \mathcal{M}) = \bigcap\{q(J') \mid J'$ is the copy of a maximal subset of I such that $J' \models \Sigma_{\text{t}}\}$. Let J be the copy of I. Then XR-certain$(q, I, \mathcal{M}) = \bigcap\{q(J') \mid J'$ is a maximal subset of J such that $J' \models \Sigma_{\text{t}}\}$, which is precisely subset-CQA$(q, J, \Sigma_{\text{t}})$. □

Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{\text{st}}, \Sigma_{\text{t}})$ be a schema mapping and q a Boolean query over $\mathbf{T}$. We consider two natural decision problems in the exchange-repair framework, and give upper bounds for their computational complexity.

- **Source-Repair Checking**: Given a source instance I and a source instance $I' \subseteq I$, is I' a source-repair of I w.r.t. $\mathcal{M}$?
- **XR-certain Query Answering**: Given a source instance I, does XR-certain$(q, I, \mathcal{M})$ evaluate to true? In other words, is $q(J')$ true on every target instance J' for which there is a source instance I' such that (I', J') is an XR-solution for I?

Theorem 3.6. *Let $\mathcal{M}$ be a* GLAV+(WAGLAV, EGD) *schema mapping.*

1. *The source-repair checking problem is in* PTIME.
2. *Let q be a union of conjunctive queries over the target schema. The* XR-certain *query answering problem for q is in* coNP.

Moreover, there is a schema mapping specified by copy s-t tgds and target egds, and a Boolean conjunctive query for which the XR-certain *query answering problem is* coNP*-complete. Thus, the data complexity of the* XR-certain *answers for Boolean conjunctive queries is* coNP*-complete.*

Proof. For the first part, the following is a polynomial time algorithm to check if $I' \subseteq I$ is a source repair of I w.r.t. $\mathcal{M}$:

> Use the chase procedure to check that I' has a solution w.r.t. $\mathcal{M}$ [19]. For every tuple $t \in I \setminus I'$, use the chase procedure to check that $I' \cup \{t\}$ does *not* have a solution w.r.t. $\mathcal{M}$.

The first step ensures that I' has a solution, and by Lemma 3.1, the second step is sufficient to ensure that I' is a maximal such subset of I. Since $\mathcal{M}$ is weakly acyclic, this algorithm runs in time which is polynomial in the size of I.

For the second part, the following is an algorithm in NP to check if XR-certain$(q, I, \mathcal{M})$ is false:

> Let I' be an arbitrary subset of I. Using the algorithm from the first part, check that I' is a source repair of I. If so, check that $q(chase(I')) = false$.

For the matching lower bound, consider the schema mapping $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{\text{st}}, \Sigma_{\text{t}})$ and target conjunctive query q, where $\mathbf{S} = \{\mathrm{P}(x,y), \mathrm{Q}(x,y)\}$, $\mathbf{T} = \{\mathrm{P}'(x,y), \mathrm{Q}'(x,y)\}$, $\Sigma_{\text{st}} = \{\mathrm{P}(x,y) \rightarrow \mathrm{P}'(x,y), \mathrm{Q}(x,y) \rightarrow \mathrm{Q}'(x,y)\}$, and $\Sigma_{\text{t}} = \{\mathrm{P}'(x,y) \wedge \mathrm{P}'(x,y') \rightarrow y = y', \mathrm{Q}'(x,y) \wedge \mathrm{Q}'(x,y') \rightarrow y = y'\}$, and $q = \exists x \exists y \exists x' \mathrm{P}'(x,y) \wedge \mathrm{Q}'(x',y)$. Note that Σ_{st} is the *copy* mapping, therefore, we have XR-certain$(q, I, \mathcal{M}) = \oplus$-CQA$(q, J, \Sigma_{\text{t}})$ where J is merely a copy of I. For the given target query and target constraints, the latter is known to be coNP-hard in data complexity [16, 21]. □

Theorem 3.6 implies that the algorithmic properties of exchange-repairs are quite different from those of $\oplus$-repairs. Indeed, as shown in [1, 18], for GLAV+(WAGLAV, EGD) schema mappings, the $\oplus$-repair-checking problem is in coNP (and can be coNP-complete), and the data complexity of the consistent answers of Boolean conjunctive queries is Π_2^p-complete [16]. This drop in complexity can be directly attributed to Theorem 3.2.

3.2 Related Work

The work reported here builds directly on the work of many others, in particular the foundational work on database repairs and consistent query answering by Arenas, Bertossi, and Chomicki [4], and on data exchange and certain query answering by Fagin et al. [19].

As mentioned earlier, the main conceptual contribution of this paper is the introduction of an inconsistency-tolerant semantics for data exchange, called exchange-repairs, in which we consider repairs to the source instance. Inconsistency-tolerant semantics have been studied in several different areas of database management, including data integration and ontology-based data access (OBDA). The common motivation for inconsistency-tolerant semantics is to give non-trivial and, in fact, meaningful semantics to query answering. We now discuss the relationship between the XR-certain answers and the inconsistency-tolerant semantics of queries in these different contexts.

3.3 Connections with data integration

In [14] and [28], the authors introduce and study the notion of *loosely-sound* semantics for queries in a data integration setting. There are two main differences between that setting and ours. To begin with, they consider schema mappings in which the schema mapping consists of GAV (global-as-view) constraints between the source (local) schema and the target (global) schema, and also of key constrains and inclusion dependencies on the target schema; in contrast, we consider richer constraint languages, namely, GLAV (global-and-local-as-view) constraints between source and target, and also target egds and target tgds. More importantly perhaps, the loosely-sound semantics are, in general, different from the XR-certain answers semantics. Specifically, given a source instance I, the loosely-sound semantics are obtained by first computing the result J of the chase of I with the GAV constraints between the source and the target, and then considering as "repairs" all instances J' that satisfy the target constraints and are inclusion maximal in their intersection with J.

If all target constraints are egds (in particular, if all target constraints are key constraints), then it is easy to show that, for target conjunctive queries, the loosely-sound semantics coincide with the consistent answers of queries with respect to subset repairs of J. Thus, in this case, the loosely-sound semantics give the same unsatisfactory answers as the materialize-then-repair approach seen in Figure 3. Concretely, this approach yields the instance J' in Figure 3 as one possible "repair" of the instance J in Figure 1, and includes the undesirable answers $(\text{peter}, \text{portman})$ and $(\text{peter}, \text{lumbergh})$ to the query $\texttt{boss}(\text{peter}, b)$. Thus, this same example shows that the loosely-sound semantics are different from the XR-certain semantics.

In [13], Calì, Lembo, and Rosati consider the notions of *loosely-sound*, *loosely-complete*, and *loosely-exact* semantics of queries on an inconsistent database. We note that the loosely-exact semantics coincide with the consistent-answer semantics with respect to symmetric-difference-repairs of the inconsistent database.

3.4 Connections with ontology-based data access

Ontology-based data access (OBDA), originally introduced in [15], is a framework for answering queries over knowledge bases. In that framework, a knowledge base over a schema $\mathbf{T}$ is a pair (D, Σ), where D is a $\mathbf{T}$-instance and Σ is a set of constraints expressed in some logical formalism over $\mathbf{T}$. The instance D represents extensional knowledge given by the facts of D, and is called the ABox. The set Σ of constraints represents intensional knowledge, and is called the TBox. In most scenarios, the schema $\mathbf{T}$ consists of unary relation symbols, called *concepts*, and of binary relation symbols, called *roles*. Moreover, Σ typically consists of sentences in some description logic. An inconsistency-tolerant semantics in the context of OBDA was first investigated in [30]; this semantics is based on the notion of *AR-repairs* (ABox-repairs) and has become known as *AR-semantics*. Subsequent investigations of AR-semantics were carried out in a number of papers, including (in chronological order) [29, 37, 9, 11, 10, 33]. These papers have analyzed the computational complexity of consistent query answering in OBDA and have also considered several variants of the AR-semantics in the OBDA framework.

Data exchange and OBDA are different frameworks that aim to formalize different aspects of data interoperability. n data exchange there are two schemas, the source schema and the target schema, with no restrictions on the type of relation symbols they contain, while in OBDA there is a single schema that typically contains only unary and binary relation symbols. Moreover, as seen in the preceding discussion, the constraints typically used in data exchange are quite different from those typically used in OBDA. One notable exception to this is the work reported in [33], where the OBDA framework studied allows for tuple-generating dependencies (it also allows for *negative constraints*, but not for equality-generating dependencies). In spite of these differences, it turns out that there are close connections between data exchange and OBDA. In what follows, we spell out these connections in detail and show that, as regards consistent query answering, each of these two frameworks can simulate the other.

We first introduce some basic concepts and terminology for OBDA; for the most part, we follow [33].

Let $\mathbf{T}$ be a schema and let (D, Σ) be a knowledge base over $\mathbf{T}$. A *model* of (D, Σ) is a $\mathbf{T}$-instance J such that $D \subseteq J$ and $J \models \Sigma$. We write $\text{mod}(D, \Sigma)$ for the set of all models of (D, Σ).

An *AR-repair* of (D, Σ) is a $\mathbf{T}$-instance D' with the following properties: (i) $D' \subseteq D$; (ii) $\text{mod}(D', \Sigma) \neq \emptyset$; (iii) D' is an inclusion maximal sub-instance of D having the second property, i.e., there is no $\mathbf{T}$-instance D'' such that $D' \subset D'' \subseteq D$ and $\text{mod}(D'', \Sigma) \neq \emptyset$. We write $\text{drep}(D, \Sigma)$ for the set of all AR-repairs of (D, Σ).

Next, we introduce the notion of consistent query answering in the context of OBDA. Let q be a Boolean query over the schema $\mathbf{T}$. We say that q is *entailed by* (D, Σ) *under AR-semantics* if for every AR-repair D' in $\text{drep}(D, \Sigma)$ and every $\mathbf{T}$-instance $J \in \text{mod}(D', \Sigma)$, we have that $J \models q$. If q is a non-Boolean query of arity k over the schema $\mathbf{T}$ and $\mathbf{a}$ is a k-tuple of constants, then we say that $q(\mathbf{a})$ is *entailed by* (D, Σ) *under AR-semantics* if $q(\mathbf{a})$ is entailed when viewed as a Boolean query; this means

that for every AR-repair D' in $\mathrm{drep}(D, \Sigma)$ and every $\mathbf{T}$-instance $J \in \mathrm{mod}(D', \Sigma)$, we have that $\mathbf{a}$ belongs to the result $q(J)$ of evaluating q on J. We write $\text{AR-certain}(q, D, \Sigma)$ to denote the set of all tuples $\mathbf{a}$ such that $q(\mathbf{a})$ is entailed by (D, Σ) under AR-semantics. By unraveling the definitions, we see that

$$\text{AR-certain}(q, D, \Sigma) = \bigcap \{q(J) : J \in \mathrm{mod}(D', \Sigma) \text{ and } D' \in \mathrm{drep}(D, \Sigma)\}.$$

We are now ready to establish the precise connections between the exchange-repairs framework and the OBDA framework.

3.4.1 From OBDA to exchange repairs

Assume that (D, Σ) is a knowledge base over a schema $\mathbf{T}$. Let $\mathbf{S}^*$ be the schema of the relation symbols occurring in D; note that $\mathbf{S}^*$ is a (possibly proper) subschema of $\mathbf{T}$. Let $\mathbf{S}$ be a *copy* of $\mathbf{S}^*$, that is, for every relation symbol R^* in $\mathbf{S}^*$, there is a relation symbol R in $\mathbf{S}$ of the same arity. If K is an $\mathbf{S}^*$-instance, we will write $K_{\mathbf{S}}$ to denote $\mathbf{S}$-copy of K, i.e., the $\mathbf{S}$-instance obtained from K by renaming the facts of K using the corresponding relation symbols in $\mathbf{S}$. Conversely, if I is an $\mathbf{S}$-instance, then we will write $I_{\mathbf{S}^*}$ to denote the $\mathbf{S}^*$-copy of I.

The next proposition tells that the OBDA framework can be simulated by the exchange-repairs framework. The proof is straightforward, and it is omitted.

Proposition 3.7. *Consider the schema mapping $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{st}, \Sigma)$, where Σ_{st} is the set of copy s-t tgds from $\mathbf{S}$ to $\mathbf{S}^*$. The following statements are true.*

1. *For every $\mathbf{T}$-instance $D' \subseteq D$ and every $\mathbf{T}$-instance J, we have that $D' \subseteq J$ and $J \models \Sigma$ if and only if J is a solution for $D'_{\mathbf{S}}$ w.r.t. $\mathcal{M}$.*
2. *For every $\mathbf{S}$-instance $I' \subseteq D_{\mathbf{S}}$ and every $\mathbf{T}$-instance J, we have that J is a solution for I' w.r.t. $\mathcal{M}$ if and only if $I'_{\mathbf{S}^*} \subseteq J$ and $J \models \Sigma$.*
3. *There is a 1-1 correspondence between the AR-repairs of (D, Σ) and the subset source repairs of $D_{\mathbf{S}}$ w.r.t. $\mathcal{M}$. In fact, they are the same up to renaming relation symbols in $\mathbf{S}^*$ by their copies in $\mathbf{S}$.*
4. *For every query q over $\mathbf{T}$, we have that $\text{AR-certain}(q, D, \Sigma) = \text{XR-certain}(q, D_{\mathbf{S}}, \mathcal{M})$.*

3.4.2 From exchange repairs to OBDA

Assume that $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{\mathrm{st}}, \Sigma_{\mathrm{t}})$ is a schema mappings in which Σ_{st} is a set of s-t tgds and Σ_{t} is a set of arbitrary constraints over $\mathbf{T}$. Recall that the schemas $\mathbf{S}$ and $\mathbf{T}$ have no relation symbols in common. The next proposition tells that the exchange-repairs framework can be simulated by the OBDA framework. Since Σ_t are arbitrary constraints, Theorem 3.2 does not necessarily apply, so we explicitly focus on *subset* source repairs.

Proposition 3.8. *Let I be a source instance. Consider the knowledge base $(I, \Sigma_{st} \cup \Sigma_t)$ with I as the ABox and the union $\Sigma_{st} \cup \Sigma_t$ over the schema $\mathbf{S} \cup \mathbf{T}$ as the TBox. The following statements are true.*

1. *For every source instance I', we have that I' is a subset source repair of I w.r.t. $\mathcal{M}$ if and only if I' is an AR-repair of $(I, \Sigma_{st} \cup \Sigma_t)$.*
2. *For every query q over $\mathbf{T}$, we have that $\text{XR-certain}(q, I, \mathcal{M}) = \text{AR-certain}(q, I, \Sigma_{st} \cup \Sigma_t)$.*

Proof. Assume first that I' is a subset source repair of I w.r.t. $\mathcal{M}$. We have to show that I' is an AR-repair of $(I, \Sigma_{\mathrm{st}} \cup \Sigma_{\mathrm{t}})$. Since I' is a subset source repair of I w.r.t. $\mathcal{M}$, we have that $I' \subseteq I$. Moreover, there is a solution J for I' w.r.t. $\mathcal{M}$. Let $J' = I' \cup J$. Clearly, we have that (i) $I' \subseteq J'$ and (ii) $J' \models \Sigma_{\mathrm{st}} \cup \Sigma_{\mathrm{t}}$, hence $\mathrm{mod}(I', \Sigma_{\mathrm{st}} \cup \Sigma_{\mathrm{t}}) \neq \emptyset$. It remains to show that I' is a maximal sub-instance of I with the preceding properties (i) and (ii). Towards a contradiction, suppose that there is a sub-instance I'' of I such that $I' \subset I'' \subseteq I$ and $\mathrm{mod}(I'', \Sigma_{\mathrm{st}} \cup \Sigma_{\mathrm{t}}) \neq \emptyset$. Consider an instance $J'' \in \mathrm{mod}(I'', \Sigma_{\mathrm{st}} \cup \Sigma_{\mathrm{t}})$. Then $I'' \subseteq J''$ and $J'' \models \Sigma_{\mathrm{st}} \cup \Sigma_{\mathrm{t}}$. Let $J''\big|_{\mathbf{T}}$ be the restriction of J'' to the target schema $\mathbf{T}$, that is, $J''\big|_{\mathbf{T}}$ is the sub-instance of J'' consisting of the facts of J'' that involve relation symbols in $\mathbf{T}$ only. We claim that $J''\big|_{\mathbf{T}}$ is a solution for I'' w.r.t. $\mathcal{M}$. Indeed, $J''\big|_{\mathbf{T}} \models \Sigma_{\mathrm{t}}$, since all formulas in Σ_{t} contain atomic formulas from $\mathbf{T}$ only. Moreover, since $I'' \subseteq J''$ and since $J'' \models \Sigma_{\mathrm{st}}$, we have $(I'', J''\big|_{\mathbf{T}}) \models \Sigma_{\mathrm{st}}$. This is so because, since Σ_{st} consists of s-t tgds, the $\mathbf{S}$-facts in $J'' \setminus I''$ play no

role in satisfying $\Sigma_{\rm st}$; we note that this may not hold if, say, $\Sigma_{\rm st}$ contained target-to-source tgds. It follows that I' is not a subset source repair for I w.r.t. $\mathcal{M}$, which is a contradiction.

Next, assume that I' is an AR-repair of $(I, \Sigma_{\rm st} \cup \Sigma_{\rm t})$. We have to show that I' is a subset source repair of I w.r.t. $\mathcal{M}$. Since I' is an AR-repair of $(I, \Sigma_{\rm st} \cup \Sigma_{\rm t})$, we have that $I' \subseteq I$ and $\text{mod}(I', \Sigma_{\rm st} \cup \Sigma_{\rm t}) \neq \emptyset$. Let J' be a member of $\text{mod}(I', \Sigma_{\rm st} \cup \Sigma_{\rm t})$. Hence, $I' \subseteq J'$ and $J' \models \Sigma_{\rm st} \cup \Sigma_{\rm t}$. If $J'\big|_{\mathbf{T}}$ is the restriction of J' to the target schema $\mathbf{T}$, then $(I', J'\big|_{\mathbf{T}}) \models \Sigma_{\rm st}$ (because $\Sigma_{\rm st}$ consists of s-t tgds) and $J'\big|_{\mathbf{T}} \models \Sigma_{\rm t}$. Thus, there is a solution for I' w.r.t. $\mathcal{M}$. Moreover, we claim that I' is a maximal sub-instance of I for which there exists a solution w.r.t. $\mathcal{M}$. Indeed, if I'' is such that $I' \subset I'' \subseteq I$ and a solution J'' for I'' w.r.t. $\mathcal{M}$ exists, then $I'' \cup J'' \models \Sigma_{\rm st} \cup \Sigma_{\rm t}$. It follows that I' is an AR-repair of $(I, \Sigma_{\rm st} \cup \Sigma_{\rm t})$, which is a contradiction.

Finally, if q is a query over $\mathbf{T}$, then, using the first part of the proposition, it is easy to show that $\text{XR-certain}(q, I, \mathcal{M}) = \text{AR-certain}(q, I, \Sigma_{\rm st} \cup \Sigma_{\rm t})$. □

4 CQA-Rewritability

In this section, we show that, for GAV+EGD schema mappings $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{\rm st}, \Sigma_{\rm t})$, it is possible to construct a set of egds Σ_s over $\mathbf{S}$ such that an $\mathbf{S}$-instance I is consistent with Σ_s if and only if I has a solution w.r.t. $\mathcal{M}$. We use this to show that $\text{XR-certain}(q, I, \mathcal{M})$ for a conjunctive query q coincides with $\text{subset-CQA}(q_s, I, \Sigma_s)$ for a union of conjunctive queries q_s. Thus, we can employ tools for consistent query answering with respect to egds in order to compute XR-certain answers for GAV+EGD schema mappings.

We will use the well-known technique of *GAV unfolding* (see, e.g., [31]). Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{\rm st}, \Sigma_{\rm t})$ be a GAV+EGD schema mapping. For each k-ary target relation $T \in \mathbf{T}$, let $q_{\rm T}$ be the set of all conjunctive queries $q(x_1, \ldots, x_k) = \exists \mathbf{y}(\phi(\mathbf{y}) \wedge x_1 = y_{i_1} \wedge \cdots \wedge x_k = y_{i_k})$, for $\phi(\mathbf{y}) \rightarrow T(y_{i_1}, \ldots, y_{i_k})$ a GAV tgd belonging to Σ_{st} (recall that we frequently omit universal quantifiers in our notation, for the sake of readability).

A *GAV unfolding* of a conjunctive query $q(\mathbf{z})$ over $\mathbf{T}$ w.r.t. Σ_{st} is a conjunctive query over $\mathbf{S}$ obtained by replacing each occurrence of a target atom $T(\mathbf{z}')$ in $q(\mathbf{z})$ with one of the conjunctive queries in $q_{\rm T}$ (substituting variables from $\mathbf{z}'$ for $x_1, \ldots, x_k$, and pulling existential quantifiers out to the front of the formula).

Similarly, we define a *GAV unfolding* of an egd $\phi(\mathbf{x}) \rightarrow x_k = x_l$ over $\mathbf{T}$ w.r.t. Σ_{st} to be an egd over $\mathbf{S}$ obtained by replacing each occurrence of a target atom $T(\mathbf{z}')$ in $\phi(\mathbf{x})$ by one of the conjunctive queries in $q_{\rm T}$ (substituting variables from $\mathbf{z}'$ for $x_1, \ldots, x_k$, and pulling existential quantifiers out to the front of the formula as needed, where they become universal quantifiers).

Figure 7 shows the GAV unfolding of the schema mapping and query from Figure 1.

Theorem 4.1. *Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{st}, \Sigma_t)$ be a* GAV+EGD *schema mapping, and let Σ_s be the set of all GAV unfoldings of egds in Σ_t w.r.t. Σ_{st}. Let I be an $\mathbf{S}$-instance. The the following are equivalent:*

1. *I satisfies Σ_s if and only if I has a solution w.r.t. $\mathcal{M}$.*
2. *The subset-repairs of I w.r.t. Σ_s are the source repairs of I w.r.t. $\mathcal{M}$.*
3. *For each conjunctive query q over $\mathbf{T}$, we have that $\text{XR-certain}(q, I, \mathcal{M}) = \textit{subset}\text{-CQA}(q_s, I, \Sigma_s)$, where q_s is the union of GAV-unfoldings of q w.r.t. Σ_{st}.*

Proof. 1. Let I be an $\mathbf{S}$-instance which does not satisfy Σ_s. Then there is an egd $\phi_1(\mathbf{x}) \wedge ... \wedge \phi_k(\mathbf{x}) \rightarrow x_i = x_j \in \Sigma_{\rm s}$ which is violated in I by some image $\phi_1(\mathbf{a}) \wedge ... \wedge \phi_k(\mathbf{a})$. By the definition of $\Sigma_{\rm s}$, there is an egd $T_1(\mathbf{x}) \wedge ... \wedge T_k(\mathbf{x}) \rightarrow x_i = x_j$ in $\Sigma_{\rm t}$, and tgds $\phi_1(\mathbf{x}) \rightarrow T_1(\mathbf{x}), ..., \phi_k(\mathbf{x}) \rightarrow T_k(\mathbf{x})$ in $\Sigma_{\rm st}$. Then for any instance J where (I, J) together satisfy $\Sigma_{\rm st}$, it holds that J contains the image $T_1(\mathbf{a}) \wedge ... \wedge T_k(\mathbf{a})$ and therefore violates $\Sigma_{\rm t}$. The proof of the converse is similar.

2. Consider that the source repairs are the maximal subsets of I for which solutions exist. Using the above, we have that these are also the maximal subsets of I which satisfy $\Sigma_{\rm s}$, and therefore they are also the subset repairs of I w.r.t. $\Sigma_{\rm s}$.

3. By definition $\text{XR-certain}(q, I, \mathcal{M})$ is the intersection over $q(J')$ for all XR-solutions (I', J') w.r.t.

$$\Sigma_s = \{\ \texttt{Task_Assignments}(p,t,d) \wedge \texttt{Task_Assignments}(p,t',d') \rightarrow d = d'\ \}$$

$$\texttt{boss}_s(person, stakeholder) = \exists task, department(\texttt{Task_Assignments}(person, task, department) \wedge \texttt{Stakeholders_old}(task, stakeholder))$$

Figure 7: The GAV Unfolding of the schema mapping and query given in Figure 1.

$\mathcal{M}$ (or in other words, for all source repairs I' and solutions J' for I' w.r.t. $\mathcal{M}$). Observe that this is the intersection of $\text{certain}(q, I', \mathcal{M})$ over all source repairs I' w.r.t. $\mathcal{M}$. We will now show that $\text{certain}(q, I', \mathcal{M}) = q_s(I')$:

> Let J' be the solution for I' w.r.t. $\mathcal{M}$. Suppose $\mathbf{a}$ is a tuple in $q(J')$. Then there is some image $T_1(\mathbf{a},\mathbf{b}) \wedge ... \wedge T_k(\mathbf{a},\mathbf{b})$ of q in J', and there are some tgds $\phi_1(\mathbf{x},\mathbf{y}) \rightarrow T_1(\mathbf{x},\mathbf{y}), ..., \phi_k(\mathbf{x},\mathbf{y}) \rightarrow T_k(\mathbf{x},\mathbf{y})$ in Σ_{st} where the image $\phi_1(\mathbf{a},\mathbf{b}) \wedge ... \wedge \phi_k(\mathbf{a},\mathbf{b})$ is in I'. By definition the clause $\exists y \phi_1(\mathbf{x},\mathbf{y}) \wedge ... \wedge \phi_k(\mathbf{x},\mathbf{y})$ is in q_s, so $\mathbf{a}$ is in $q_s(I')$. The proof of the converse is similar.

We now have that $\text{XR-certain}(q, I, \mathcal{M})$ is the intersection over $q_s(I')$ for all source repairs I' of I w.r.t. $\mathcal{M}$. By the second item of the theorem, this gives the intersection over $q_s(I')$ for all subset repairs I' of I w.r.t. Σ_{s}, which is simply $\text{subset-CQA}(q_s, I, \Sigma_{\text{s}})$.

□

The following result tells us that Theorem 4.1 cannot be extended to schema mappings containing LAV s-t tgds.

Theorem 4.2. *Consider the* LAV+EGD *schema mapping* $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{st}, \Sigma_t)$*, where*

- $\mathbf{S} = \{R\}$ *and* $\mathbf{T} = \{T\}$*,*
- $\Sigma_{st} = \{R(x,y) \rightarrow \exists u\ T(x,u) \wedge T(y,u)\}$*, and*
- $\Sigma_t = \{T(x,y) \wedge T(x,z) \rightarrow y = z\}$*.*

Consider the query $q(x,y) = \exists z.\ T(x,z) \wedge T(y,z)$ *over* $\mathbf{T}$*. There does not exist a UCQ* q_s *over* $\mathbf{S}$ *and a set of universal first-order sentences (in particular, egds)* Σ_s *such that, for every instance* I*, we have that* $\text{XR-certain}(q, I, \mathcal{M}) = \textit{subset}\text{-CQA}(q_s, I, \Sigma_s)$*.*

It is worth noting that the schema mapping $\mathcal{M}$ in the statement of Theorem 4.2 is such that every source instance has a solution, and hence "XR-certain" could be replaced by "certain" in the statement.

Proof. We start by observing that $\text{certain}(q, I, \mathcal{M})$ expresses undirected reachability along the relation R:

Claim: For every $\mathbf{S}$-instance I, $\text{certain}(q, I, \mathcal{M}) = \{(a,b) \in adom(I) \mid b \text{ is reachable from } a \text{ by an undirected } R\text{-path}\}$.

The left-to-right inclusion can be proved by induction on the length of the shortest undirected path from a to b, while, for the right-to-left inclusion, it is enough to consider the solution J that contains a null value for each connected component of I, and such that J contains all facts of the form $T(a, N)$ for $a \in adom(I)$, where N is the null value associated to the connected component of I to which a belongs.

Now, suppose for the sake of a contradiction that q_s and Σ_s as described in the statement of the proposition exist. Let k be the number of variables in q_s. Let I be an instance that consists of a directed path of length $k+1$ from a to b. It follows from the above claim, and from our assumption on q_s and Σ_s, that $(a,b) \in \text{subset-CQA}(q_s, I, \Sigma_s)$, and for every proper subinstance I' of I, we have that $(a,b) \notin \text{certain}(q, I', M)$.

Claim: The instance I is consistent with Σ_s.

Suppose for the sake of a contradiction that the above claim does not hold. Let I' be any subset-repair of I with respect to Σ_s. Since I' is a

proper sub-instance of I, we have that $(a,b) \notin \text{certain}(q, I', \Sigma)$. In particular, since I' satisfies Σ_s, we have that $(a,b) \notin q_s(I')$. But since I' is a repair of I, this means that $(a,b) \notin \text{subset-CQA}(q_s, I, \Sigma_s)$, a contradiction.

Since $(a,b) \in \text{subset-CQA}(q_s, I, \Sigma_s)$ and I is consistent with Σ_s we have that $(a,b) \in q_s(I)$. That is, there is a homomorphism h from q_s to I. Let I'' be the sub-instance of I consisting of the facts involving only values that are in the image of h. Since I contains k facts and q contains $k+1$ facts, I'' is a proper sub-instance of I. Moreover, since universal first-order sentences are preserved under taking induced sub-instances, every egd true in I is also true in I'' and therefore, I'' is consistent with Σ_s. Finally, by construction, $q_s(I'') = \mathit{true}$. Therefore, $(a,b) \in \text{subset-CQA}(q_s, I'', \Sigma_s)$. This contradicts the fact that $(a,b) \notin \text{certain}(q, I'', M)$.

□

The following result tells us that Theorem 4.1 also cannot be extended to schema mappings containing GAV target tgds.

Theorem 4.3. *Consider the* GAV+(GAV, EGD) *schema mapping* $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{st}, \Sigma_t)$*, where*

- $\mathbf{S} = \{R\}$ *and* $\mathbf{T} = \{T\}$*,*
- $\Sigma_{st} = \{R(x,y) \rightarrow T(x,y)\}$*, and*
- $\Sigma_t = \{T(x,y) \wedge T(y,z) \rightarrow T(x,z)\}$*.*

Consider the query $q(x,y) = T(x,y)$ *over* $\mathbf{T}$*. There does not exist a UCQ* q_s *over* $\mathbf{S}$ *and a set of universal first-order sentences (in particular, egds)* Σ_s *such that, for every instance* I*, we have that* $\text{XR-certain}(q, I, \mathcal{M}) = \textit{subset}\text{-CQA}(q_s, I, \Sigma_s)$*.*

Proof. We start by observing that $\text{certain}(q, I, \mathcal{M})$ expresses directed reachability along the relation R: for every $\mathbf{S}$-instance I, $\text{certain}(q, I, \mathcal{M}) = \{(a,b) \in \mathit{adom}(I) \mid b \text{ is reachable from } a \text{ by a directed } R\text{-path}\}$. The claim is proved by induction on the length of the path. The remainder of the proof is identical to that of Theorem 4.2 (the difference between directed paths and undirected paths is inessential to the argument).

□

5 DLP-Rewritability

We saw in the previous section that the applicability of the CQA-rewriting approach is limited to GAV+EGD schema mappings. In this section, we consider another approach to computing XR-certain answers, based on a reduction to the problem of computing certain answers over the stable models of a disjunctive logic program. Our reduction is applicable to GLAV+(WAGLAV, EGD) schema mappings. First, we reduce the case of GLAV+(WAGLAV, EGD) schema mappings to the case of GAV+(GAV, EGD) schema mappings.

Theorem 5.1. *From a* GLAV+(WAGLAV, EGD) *schema mapping* $\mathcal{M}$ *we can construct a* GAV+(GAV, EGD) *schema mapping* $\hat{\mathcal{M}}$ *such that, from a conjunctive query* q*, we can construct a union of conjunctive queries* $\hat{q}$ *with* $\text{XR-certain}(q, I, \mathcal{M}) = \text{XR-certain}(\hat{q}, I, \hat{\mathcal{M}})$*.*

The proof of Theorem 5.1 is given in Section 6 (it is entailed by Theorem 6.2). Theorem 4.3 shows that the CQA-rewriting approach studied in Section 4 is, in general, not applicable to GAV+(GAV, EGD) schema mappings and unions of conjunctive queries. To address this problem, we will now consider a different approach to computing XR-certain answers, using disjunctive logic programs. Although stable models are popular in the literature, including for database repairs, we find that the selective minimization offered by parallel circumscription is a better fit for XR-certain semantics because our minimality condition applies only to the source-part of the schema. We then use a result from [25] to translate back into the realm of stable models.

Stable models of disjunctive logic programs have been well-studied as a way to compute database repairs ([34] provides a thorough treatment). In [14], Calì et al. give an encoding of their loosely-sound semantics for data integration as a disjunctive logic program. Their encoding is applicable for *non-key-conflicting* sets of constraints, a syntactic condition that is orthogonal to *weak acyclicity*, and which eliminates the utility of named nulls. Although their semantics use a notion of minimality that is similar to

ours, our setting and our syntactic condition differ sufficiently that our results are complementary.

Fix a domain *Const*. A *disjunctive logic program* (DLP) Π over a schema $\mathbf{R}$ is a finite collection of rules of the form

$$\alpha_1 \vee \ldots \vee \alpha_n \leftarrow \beta_1, \ldots, \beta_m, \neg\gamma_1, \ldots, \neg\gamma_k.$$

where $n, m, k \geq 0$ and $\alpha_1, \ldots, \alpha_n, \beta_1, \ldots, \beta_m, \gamma_1, \ldots, \gamma_k$ are atoms formed from the relations in $\mathbf{R} \cup \{=\}$, using the constants in *Const* and first-order variables. A DLP is said to be *positive* if it consists of rules that do not contain negated atoms except possibly for inequalities. A DLP is said to be *ground* if it consists of rules that do not contain any first-order variables. A *model* of Π is an $\mathbf{R}$-instance I over domain *Const* that satisfies all rules of Π (viewed as universally quantified first-order sentences). A rule in which $n = 0$ is called a constraint, and is satisfied only if its body is not satisfied. A *minimal model* of Π is a model M of Π such that there does not exist a model M' of Π where the facts of M' form a strict subset of the facts of M. More generally, for subsets $\mathbf{R}_{\text{M}}, \mathbf{R}_{\text{F}} \subseteq \mathbf{R}$, an $\langle \mathbf{R}_{\text{M}}, \mathbf{R}_{\text{F}} \rangle$*-minimal model* of Π is a model M of Π such that there does not exist a model M' of Π where the facts of M' involving relations from $\mathbf{R}_{\text{M}}$ form a strict subset of the facts of M involving relations from $\mathbf{R}_{\text{M}}$, and the set of facts of M' involving relations from $\mathbf{R}_{\text{F}}$ is equal to the set of facts of M involving relations from $\mathbf{R}_{\text{F}}$ [25]. Although minimal models are a well-behaved semantics for positive DLPs, it is not well suited for programs with negations. The *stable model* semantics is a widely used semantics of DLPs that are not necessarily positive. For positive DLPs, it coincides with the minimal model semantics. For a ground DLP Π over a schema $\mathbf{R}$ and an $\mathbf{R}$-instance M over the domain *Const*, the *reduct* Π^M of Π with respect to M is the DLP containing, for each rule $\alpha_1 \vee \ldots \vee \alpha_n \leftarrow \beta_1, \ldots, \beta_m, \neg\gamma_1, \ldots, \neg\gamma_k$, with $M \not\models \gamma_i$ for all $i \leq k$, the rule $\alpha_1 \vee \ldots \vee \alpha_n \leftarrow \beta_1, \ldots, \beta_m$. A *stable model* of a ground DLP Π is an $\mathbf{R}$-instance M over the domain *Const* such that M is a minimal model of the reduct Π^M. See [22] for more details.

In this section, we will construct positive DLP programs whose $\langle \mathbf{R}_{\text{M}}, \mathbf{R}_{\text{F}} \rangle$-minimal models correspond to XR-solutions. In light of Theorem 5.1, we may restrict our attention to GAV+(GAV, EGD) schema mappings.

In [25] it was shown that a positive ground DLP Π over a schema $\mathbf{R}$, together with subsets $\mathbf{R}_{\text{M}}, \mathbf{R}_{\text{F}} \subseteq \mathbf{R}$, can be translated in polynomial time to a (not necessarily positive) DLP Π' over a possibly larger schema that includes $\mathbf{R}$, such that there is a bijection between the $\langle \mathbf{R}_{\text{M}}, \mathbf{R}_{\text{F}} \rangle$-minimal models of Π and the stable models of Π', where every pair of instances that stand in the bijection agree on all facts over the schema $\mathbf{R}$. This shows that DLP reasoners based on the stable model semantics, such as DLV [32, 2], can be used to evaluate positive ground disjunctive logic programs under the $\langle \mathbf{R}_{\text{M}}, \mathbf{R}_{\text{F}} \rangle$-minimal model semantics. Although stated only for ground programs in [25], this technique can be used for arbitrary positive DLPs through grounding. Note that, when a program is grounded, inequalities are reduced to $\top$ or $\bot$.

Theorem 5.2. *Given a* GAV+(GAV, EGD) *schema mapping* $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{st}, \Sigma_t)$*, we can construct in linear time a positive DLP* Π *over a schema* $\mathbf{R}$ *that contains* $\mathbf{S} \cup \mathbf{T}$*, and subsets* $\mathbf{R}_{\text{M}}, \mathbf{R}_{\text{F}} \subseteq \mathbf{R}$*, such that for every union* q *of conjunctive queries over* $\mathbf{T}$ *and for every* $\mathbf{S}$*-instance* I*, we have that* XR-certain$(q, I, \mathcal{M}) = \bigcap\{q(M) \mid M$ *is an* $\langle \mathbf{R}_{\text{M}}, \mathbf{R}_{\text{F}} \rangle$*-minimal model of* $\Pi \cup I\}$*.*

Proof. We construct a disjunctive logic program $\Pi_{\text{XRc}}(\mathcal{M})$ for a GAV+(GAV, EGD) schema mapping $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{\text{st}}, \Sigma_{\text{t}})$ as follows:

1. For each source relation S with arity n, add the rules

   $$S_k(x_1, \ldots, x_n) \vee S_d(x_1, \ldots, x_n) \leftarrow S(x_1, \ldots, x_n)$$
   $$\bot \leftarrow S_k(x_1, \ldots, x_n), S_d(x_1, \ldots, x_n)$$
   $$S(x_1, \ldots, x_n) \leftarrow S_k(x_1, \ldots, x_n)$$

   where S_k and S_d represent the *kept* and *deleted* atoms of S, respectively.

2. For each s-t tgd $\phi(\mathbf{x}) \rightarrow T(\mathbf{x}')$ in Σ_{st}, add the rule

   $$T(\mathbf{x}') \leftarrow \alpha_1, \ldots, \alpha_m$$

   where $\alpha_1, \ldots, \alpha_m$ are the atoms in $\phi(\mathbf{x})$, in which each relation S has been uniformly replaced by S_k.

3. For each tgd $\phi(\mathbf{x}) \rightarrow T(\mathbf{x}')$ in Σ_{t}, add the rule

$$T(\mathbf{x}) \leftarrow \alpha_1, \ldots, \alpha_m$$

where $\alpha_1, \ldots, \alpha_m$ are the atoms in $\phi(\mathbf{x})$.

4. For each egd $\phi(\mathbf{x}) \rightarrow x_1 = x_2$, where $x_1, x_2 \in \mathbf{x}$, add the rule

$$\bot \leftarrow \alpha_1, \ldots, \alpha_m, x_1 \neq x_2,$$

where $\alpha_1, \ldots, \alpha_m$ are the atoms in $\phi(\mathbf{x})$.

We minimize the model w.r.t. $\mathbf{R}_{\mathrm{M}} = \{S_d \mid S \in \mathbf{S}\}$, and fix $\mathbf{R}_{\mathrm{F}} = \{S \mid S \in \mathbf{S}\}$. The disjunctive logic program for $\mathcal{M}$, denoted $\Pi_{\mathrm{XRc}}(\mathcal{M})$, is a straightforward encoding of the constraints in Σ_{st} and Σ_{t} as disjunctive logic rules over an indefinite view of the source instance. Since the source instance is fixed, the rules of the form $S(x_1, \ldots, x_n) \leftarrow S_k(x_1, \ldots, x_n)$ in $\Pi_{\mathrm{XRc}}(\mathcal{M})$ force the kept atoms to be a sub-instance of the source instance. Notice that egds are encoded as denial constraints, and that disjunction is used only to non-deterministically choose a subset of the source instance.

To prove the theorem, we first show that the restriction of every $\langle \mathbf{R}_{\mathrm{M}}, \mathbf{R}_{\mathrm{F}} \rangle$-minimal model of $\Pi_{\mathrm{XRc}}(\mathcal{M}) \cup I$ to the schema $\{S_k \mid S \in \mathbf{S}\} \cup \mathbf{T}$ constitutes an exchange-repair solution. We then show that for every exchange-repair solution, we can build a corresponding $\langle \mathbf{R}_{\mathrm{M}}, \mathbf{R}_{\mathrm{F}} \rangle$-minimal model of $\Pi_{\mathrm{XRc}}(\mathcal{M}) \cup I$.

Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{\mathrm{st}}, \Sigma_{\mathrm{t}})$ be a GAV+(GAV, EGD) schema mapping. Let $\Pi = \Pi_{\mathrm{XRc}}(\mathcal{M})$. Let $\mathbf{R}$ be the schema of Π, and let $\mathbf{R}_{\mathrm{M}} = \{S_d \mid S \in \mathbf{S}\}$, and $\mathbf{R}_{\mathrm{F}} = \{S \mid S \in \mathbf{S}\}$. Let q be a union of conjunctive queries over $\mathbf{T}$, and let I be an $\mathbf{S}$-instance.

We first prove that a certain restriction of every $\langle \mathbf{R}_{\mathrm{M}}, \mathbf{R}_{\mathrm{F}} \rangle$-minimal model of $\Pi \cup I$ is an exchange-repair solution. Let M be an $\langle \mathbf{R}_{\mathrm{M}}, \mathbf{R}_{\mathrm{F}} \rangle$-minimal model of $\Pi \cup I$. Then for each source atom $S(c_1, \ldots, c_n)$, M satisfies $S_k(c_1, \ldots, c_n) \vee S_d(c_1, \ldots, c_n) \leftarrow S(c_1, \ldots, c_n)$ and $\leftarrow S_k(c_1, \ldots, c_n), S_d(c_1, \ldots, c_n)$ and therefore contains exactly one of $S_k(c_1, \ldots, c_n)$ or $S_d(c_1, \ldots, c_n)$. Furthermore, for every atom $S_k(d_1, \ldots, d_n) \in M$, M satisfies $S(d_1, \ldots, d_n) \leftarrow S_k(d_1, \ldots, d_n)$. Let I' be a renaming of the restriction of M to the *kept* predicates (by removal of the $_k$ subscript), and observe that since I is fixed, I' is a sub-instance of I. Furthermore, since $M \models \Pi \cup I$ (which contains copies of the constraints of $\mathcal{M}$ over its *kept* predicates), I' has a solution w.r.t. $\mathcal{M}$. Finally, an appropriate renaming (by removal of the $_d$ subscript) of the restriction of M to $\mathbf{R}_{\mathrm{M}}$ (the *deleted* predicates) is equal to $I \setminus I'$, and since M is a $\langle \mathbf{R}_{\mathrm{M}}, \mathbf{R}_{\mathrm{F}} \rangle$-minimal model of $\Pi \cup I$, we have that there is no I'' such that $I' \subset I'' \subseteq I$ and a solution exists for I'' w.r.t. $\mathcal{M}$. Therefore, I' is a source repair of I w.r.t. $\mathcal{M}$, and I' along with the restriction of M to $\mathbf{T}$ is an exchange-repair solution for I w.r.t. $\mathcal{M}$.

We now prove that for every exchange-repair solution, there exists an $\langle \mathbf{R}_{\mathrm{M}}, \mathbf{R}_{\mathrm{F}} \rangle$-minimal model of $\Pi \cup I$. Let (I', J') be an exchange-repair solution for I w.r.t $\mathcal{M}$. Let $M = I \cup I'_k \cup (I \setminus I')_d \cup J'$, where I'_k is I' renamed over the *kept* predicates, and $(I \setminus I')_d$ is $I \setminus I'$ renamed over the *deleted* predicates. Since I' is a subset of I, and $I \setminus I'$ is disjoint from I', we have that the rules of the forms $S_k(x_1, \ldots, x_n) \vee S_d(x_1, \ldots, x_n) \leftarrow S(x_1, \ldots, x_n)$, and $\leftarrow S_k(x_1, \ldots, x_n), S_d(x_1, \ldots, x_n)$, and $S(x_1, \ldots, x_n) \leftarrow S_k(x_1, \ldots, x_n)$ are satisfied. It also holds that (I', J') satisfy $\mathcal{M}$, and therefore M is a model of $\Pi \cup I$. Finally, since there is no I'' such that $I' \subset I'' \subseteq I$ and a solution exists for I'' w.r.t. $\mathcal{M}$, M is also $\langle \mathbf{R}_{\mathrm{M}}, \mathbf{R}_{\mathrm{F}} \rangle$-minimal.

Therefore, XR-certain$(q, I, \mathcal{M}) = \bigcap \{q(M) \mid M$ is an $\langle \mathbf{R}_{\mathrm{M}}, \mathbf{R}_{\mathrm{F}} \rangle$-minimal model of $\Pi \cup I\}$.

□

Figure 8 illustrates the disjunctive logic program obtained for the schema mapping from Figure 1.

6 From GLAV+(WAGLAV, EGD) to GAV+(GAV, EGD)

In this section, we prove Theorem 5.1, and discuss some additional literature related to this particular result. Let $\mathcal{M}_1$ and $\mathcal{M}_2$ be schema mappings with the same source schema. We will write $\mathcal{M}_1 \rightsquigarrow_{\mathrm{UCQ}} \mathcal{M}_2$ if for every UCQ q over the target schema of $\mathcal{M}_1$, there is a UCQ q' over the target schema of $\mathcal{M}_2$ such that for all source instance I, XR-certain$(q, I, \mathcal{M}_1) =$ XR-certain$(q', I, \mathcal{M}_2)$. Using this notation, Theorem 5.1 states that *for ev-*

$\texttt{Task_Assignments}_k(p,t,d) \vee \texttt{Task_Assignments}_d(p,t,d) \leftarrow \texttt{Task_Assignments}(p,t,d).$
$\bot \leftarrow \texttt{Task_Assignments}_k(p,t,d), \texttt{Task_Assignments}_d(p,t,d).$
$\texttt{Task_Assignments}(p,t,d) \leftarrow \texttt{Task_Assignments}_k(p,t,d).$

$\texttt{Stakeholders_old}_k(t,s) \vee \texttt{Stakeholders_old}_d(t,s) \leftarrow \texttt{Stakeholders_old}(t,s).$
$\bot \leftarrow \texttt{Stakeholders_old}_k(t,s) \wedge \texttt{Stakeholders_old}_d(t,s).$
$\texttt{Stakeholders_old}(t,s) \leftarrow \texttt{Stakeholders_old}_k(t,s).$

$\texttt{Departments}(p,d) \leftarrow \texttt{Task_Assignments}_k(p,t,d).$
$\texttt{Tasks}(p,t) \leftarrow \texttt{Task_Assignments}_k(p,t,d).$
$\texttt{Stakeholders_new}(t,s) \leftarrow \texttt{Stakeholders_old}_k(t,s).$

$\bot \leftarrow \texttt{Departments}(p,d), \texttt{Departments}(p,d'), d \neq d'.$

$\texttt{boss}(person, stakeholder) \leftarrow \texttt{Tasks}(person, task), \texttt{Stakeholders_new}(task, stakeholder)$

Figure 8: The disjunctive logic program over $\langle \mathbf{R}_{\text{M}}, \mathbf{R}_{\text{F}} \rangle$-minimal models for the schema mapping and query given in Figure 1. Here, $\mathbf{R}_{\text{M}} = \{\texttt{Task_Assignments}_d, \texttt{Stakeholders_old}_d\}$ and $\mathbf{R}_{\text{F}} = \{\texttt{Task_Assignments}, \texttt{Stakeholders_old}\}$

ery GLAV+(WAGLAV, EGD) *schema mapping* $\mathcal{M}$ *there is a* GAV+(GAV, EGD) *schema mapping* $\mathcal{M}'$ *with* $\mathcal{M} \rightsquigarrow_{\text{UCQ}} \mathcal{M}'$. We will in fact prove a stronger statement that applies to schema mappings defined by *second-order tgds*. Second-order tgds serve not only to strengthen the result, but also to make its proof more natural.

6.1 Second-order TGDs

Second-order tgds are a natural extension of tgds that was introduced in [20] in the context of schema mapping composition. We recall the definition.

Let $\mathbf{f}$ be a collection of function symbols, each having a designated arity. A *simple term* is a constant or variable. A *compound term* is a function applied to a list of terms, such that the arity of the function symbol is respected. By an $\mathbf{f}$-*term*, we mean either a simple term, or a compound term built up from variables and/or constants using the function symbols in $\mathbf{f}$. We will omit $\mathbf{f}$ from the notation when it is understood from context. The *depth* of a term is the maximal nesting of function symbols, with $depth(e) = 0$ when e is a simple term. A *ground term* is a term in which no variables appear.

A *second-order tgd* (SO tgd) over a schema $\mathbf{R}$ is an expression of the form

$$\sigma = \exists \mathbf{f}\, (\forall \mathbf{x_1}(\phi_1 \rightarrow \psi_1) \wedge \cdots \wedge \forall \mathbf{x_n}(\phi_n \rightarrow \psi_n))$$

where $\mathbf{f}$ is a collection of function symbols, and

1. each ϕ_i is a conjunction of (a) atoms $S(y_1, \ldots, y_k)$ where $S \in \mathbf{R}$ and $y_1, \ldots, y_k$ are variables from $\mathbf{x_i}$; and (b) equalities of the form $t_1 = t_2$ where t_1, t_2 are terms over $\mathbf{x_i}$ and $\mathbf{f}$.
2. each ψ_i is a conjunction of atoms $S(t_1, \ldots, t_k)$ where $S \in \mathbf{R}$ and $t_1, \ldots, t_k$ are $\mathbf{f}$-terms built from $\mathbf{x_i}$.
3. each variable in $\mathbf{x_i}$ occurs in a relational atom in ϕ_i.

We say that an $\mathbf{R}$-instance I satisfies σ if there exists a collection of functions $\mathbf{f^0}$ (whose domain and range are Const $\cup$ Nulls) such that each "clause" $\forall \mathbf{x_i}(\phi_i \rightarrow \psi_i)$ of σ is satisfied in I where each function symbol in $\mathbf{f}$ is interpreted by the corresponding function in $\mathbf{f^0}$. We will write $I \models \sigma$ when this is the case, or, if we wish to make $\mathbf{f^0}$ explicit in the notation, $I \models \sigma\ [\mathbf{f} \mapsto \mathbf{f^0}]$.

A *source-to-target* SO tgd for source schema $\mathbf{S}$ and target schema $\mathbf{T}$ is an SO tgd over $\mathbf{S} \cup \mathbf{T}$, of the above

form, where each ϕ_i contains only relation symbols from $\mathbf{S}$ and each ψ_i contains only relation symbols from $\mathbf{T}$. We note that, in [20], only source-to-target SO tgds were considered.

An *equality-free* SO tgd (EFSOTGD) is an SO tgd that does not contain term equalities. We denote by SOTGD+(SOTGD, EGD) the class of schema mappings $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{\text{st}}, \Sigma_{\text{t}})$ where Σ_{st} is a set of source-to-target SO tgds over $\mathbf{S}$ and $\mathbf{T}$, and Σ_t is a set of SO tgds and/or egds over $\mathbf{T}$. Other classes of schema mappings, such as SOTGD+SOTGD and EFSOTGD+EFSOTGD, are defined analogously. Note that an important subclass of equality-free SO tgds are the *plain* SO tgds, introduced in [6], in which no terms contain nested functions.

It is known that every tgd is logically equivalent to a SO tgd, which can be obtained from it by *skolemization* [20]. Although stated in the literature only for the case of source-to-target tgds [20], the same applies to target tgds. Figure 10 shows the skolemization of the example schema mapping in Figure 9.

Moreover, if we adapt the concept of weak acyclicity to SO tgds in the appropriate way, then every weakly acyclic set of tgds is logically equivalent to a weakly acyclic SO tgd.

More precisely, we say that a set Σ of SO tgds is weakly acyclic if there is no cycle in its dependency graph containing a special edge, where the dependency graph associated to a set of SO tgds is defined as follows:

1. the directed graph whose nodes are positions (R, i) where R is a relation symbol and i is an attribute position of R (as before)
2. there is a normal edge from (R, i) to (S, j) if Σ contains a SO tgd of the form

$$\sigma = \exists \mathbf{f}\, (\forall \mathbf{x_1}(\phi_1 \to \psi_1) \wedge \cdots \wedge \forall \mathbf{x_n}(\phi_n \to \psi_n))$$

   and for some $i \leq n$, ϕ_i contains a variable in position (R, i) and ψ_i contains the same variable in position (S, j).
3. there is a special edge from (R, i) to (S, j) if Σ contains a SO tgd of the form

$$\sigma = \exists \mathbf{f}\, (\forall \mathbf{x_1}(\phi_1 \to \psi_1) \wedge \cdots \wedge \forall \mathbf{x_n}(\phi_n \to \psi_n))$$

   and for some $i \leq n$, ϕ_i contains a variable in position (R, i) and ψ_i contains a compound term in position (S, j) containing the same variable.

We then have:

Proposition 6.1. *Every* GLAV+(WAGLAV, EGD) *schema mapping is logically equivalent to a weakly acyclic* EFSOTGD+(EFSOTGD, EGD) *schema mapping.*

Indeed, if $\mathcal{M}$ is a GLAV+(WAGLAV, EGD) schema mapping, and $\mathcal{M}'$ is the EFSOTGD+(EFSOTGD, EGD) schema mapping obtained from $\mathcal{M}$ by skolemization, then $\mathcal{M}$ and $\mathcal{M}'$ are logically equivalent. Moreover, it is easy to see that $\mathcal{M}$ and $\mathcal{M}'$ have the same dependency graph, and, therefore, $\mathcal{M}'$ is weakly acyclic.

In the remainder of this section, we will establish:

Theorem 6.2. *For every weakly acyclic* SOTGD+(SOTGD, EGD) *schema mapping $\mathcal{M}$ there is a* GAV+(GAV, EGD) *schema mapping $\mathcal{M}'$ such that $\mathcal{M} \rightsquigarrow_{\text{UCQ}} \mathcal{M}'$.*

The proof borrows ideas from previous literature, and we discuss relevant related work at the end of the section.

$$\begin{array}{l}
\mathbf{S} = \{\mathrm{R}\} \\
\mathbf{T} = \{\mathrm{T}\} \\
\Sigma_{\text{st}} = \{\ \mathrm{R}(x,y) \ \to\ \exists u(\mathrm{T}(x,u) \wedge \mathrm{T}(y,u))\ \} \\
\Sigma_{\text{t}} = \{\ \mathrm{T}(x,y) \wedge \mathrm{T}(x,y') \to y = y'\ \} \\
\mathrm{q}(x,y) := \exists u(\mathrm{T}(x,u) \wedge \mathrm{T}(y,u))
\end{array}$$

Figure 9: An example schema mapping and query.

$$\begin{array}{l}
\mathbf{S} = \{\mathrm{R}\} \\
\mathbf{T} = \{\mathrm{T}\} \\
\Sigma_{\text{st}} = \left\{ \exists f \left(\begin{array}{lll} \mathrm{R}(x,y) & \to & \mathrm{T}(x, f(x,y)) \wedge \\ \mathrm{R}(x,y) & \to & \mathrm{T}(y, f(x,y)) \end{array} \right) \right\} \\
\Sigma_{\text{t}} = \{\ \mathrm{T}(x,y) \wedge \mathrm{T}(x,y') \to y = y'\ \} \\
\mathsf{q}(x,y) := \exists u(\mathrm{T}(x,u) \wedge \mathrm{T}(y,u))
\end{array}$$

Figure 10: Result of skolemizing the schema mapping in Figure 9.

6.2 Eliminating Equalities to Establish Freeness

In this section, we will rewrite our schema mapping to eliminate egds as well as equality conditions in SO tgds. This allows us to work with solutions in which there is a one-to-one correspondence between ground terms (of any depth) and their values. This property, called *freeness*, which we define below, is used in Section 6.3.

For simplicity we first restrict attention to EFSOTGD+(SOTGD, EGD) schema mappings.

Definition 6.1 (Equality Singularization)**.** Fix a fresh binary relation symbol Eq.

- The *equality singularization* of a conjunctive query $q(\mathbf{x}) = \exists \mathbf{y}\phi(\mathbf{x},\mathbf{y})$, denoted by $q^{\mathrm{Eq}}(\mathbf{x})$, is the conjunctive query $\exists \mathbf{yz}\phi'(\mathbf{x},\mathbf{y},\mathbf{z})$ obtained from q as follows: whenever a variable u (free or quantified) occurs more than once in ϕ, we replace each occurence other than the first occurrence by a fresh distinct variable z and we add the atom $Eq(u,z)$.
- The *equality singularization* of an egd

  $$\sigma = \forall \mathbf{x}(\phi \to x_i = x_j)$$

  is the GAV tgd

  $$\sigma^{Eq} = \forall \mathbf{xz}(\phi' \to Eq(x_i, x_j))$$

  where $\phi^{Eq} = \exists \mathbf{z}\phi'$
- The *equality singularization* of an SO tgd

  $$\sigma = \exists \mathbf{f} \bigwedge_{i=1\ldots n} (\forall \mathbf{x}_i(\phi_i \wedge \alpha_i \to \psi_i))$$

  (where each ϕ_i is a conjunction of relational atoms and each α_i is a conjunction of equalities) is the equality-free SO tgd

  $$\sigma' = \exists \mathbf{f} \bigwedge_{i=1\ldots n} (\forall \mathbf{x}_i\mathbf{z}_i(\phi'_1 \wedge \alpha'_i \to \psi_i))$$

  where $\phi_i^{Eq} = \exists \mathbf{z}_i\phi'$ and α'_i is obtained from α_i by replacing each equality $s = t$ by $Eq(s,t)$.
- The *equality singularization* of a EFSOTGD+(SOTGD, EGD) schema mapping $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{\mathrm{st}}, \Sigma_{\mathrm{t}})$ is the EFSOTGD+EFSOTGD schema mapping

  $$\mathcal{M}^{\mathrm{Eq}} = (\mathbf{S}, \mathbf{T} \cup \{\mathrm{Eq}\}, \Sigma_{\mathrm{st}}, \{\sigma^{\mathrm{Eq}} \mid \sigma \in \Sigma_{\mathrm{t}}\} \cup \mathrm{eqAx}(\mathbf{T}))$$

  where $\mathrm{eqAx}(\mathbf{T})$ is the set of (full) tgds of the form $\mathrm{T}(x_1,\ldots,x_n) \to \mathrm{Eq}(x_1,x_1) \wedge \cdots \wedge \mathrm{Eq}(x_n,x_n)$ where T is a relation in $\mathbf{T}$, along with the tgds $\mathrm{Eq}(x_1,x_2) \to \mathrm{Eq}(x_2,x_1)$ and $\mathrm{Eq}(x_1,x_2) \wedge \mathrm{Eq}(x_2,x_3) \to \mathrm{Eq}(x_1,x_3)$.

$$\begin{array}{ll}
\mathbf{S} & = \{\mathrm{R}\} \\
\mathbf{T} & = \{\mathrm{T}, \mathrm{Eq}\} \\
\Sigma_{\mathrm{st}} & = \left\{ \exists f \left(\begin{array}{lll} \mathrm{R}(x,y) & \to & \mathrm{T}(x, f(x,y)) \wedge \\ \mathrm{R}(x,y) & \to & \mathrm{T}(y, f(x,y)) \end{array} \right) \right\} \\
\Sigma_{\mathrm{t}}^{\mathrm{Eq}} & = \left\{ \begin{array}{l} \mathrm{T}(x,y) \wedge \mathrm{Eq}(x,x') \wedge \mathrm{T}(x',y') \to \mathrm{Eq}(y,y') \\ \mathrm{T}(x,y) \to \mathrm{Eq}(x,x) \\ \mathrm{T}(x,y) \to \mathrm{Eq}(y,y) \\ \mathrm{Eq}(x_1,x_2) \to \mathrm{Eq}(x_2,x_1) \\ \mathrm{Eq}(x_1,x_2) \wedge \mathrm{Eq}(x_2,x_3) \to \mathrm{Eq}(x_1,x_3) \end{array} \right\} \\
\mathrm{q}^{\mathrm{Eq}}(x,y) & := \exists u,u'(\mathrm{T}(x,u) \wedge \mathrm{Eq}(u,u') \wedge \mathrm{T}(y,u'))
\end{array}$$

Figure 11: Equality singularization of the schema mapping and query from Figure 10.

Figure 11 shows equality singularization in action.

Proposition 6.3. *Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{st}, \Sigma_t)$ be an* SOTGD+(SOTGD, EGD) *schema mapping and let $\mathcal{M}^{\mathrm{Eq}}$ be its equality singularization. For every $\mathbf{S}$-instance I,*

1. *I has a solution w.r.t. $\mathcal{M}$ if and only if I has a solution J' w.r.t. $\mathcal{M}^{\mathrm{Eq}}$ such that there is no pair of distinct constants a, b where $J' \models \mathrm{Eq}(a,b)$.*
2. *If I has a solution w.r.t. $\mathcal{M}$, then for every UCQ q over $\mathbf{T}$, $\mathrm{certain}(q, I, \mathcal{M}) = \mathrm{certain}(q^{\mathrm{Eq}}, I, \mathcal{M}^{\mathrm{Eq}})$.*

Proof. [⇒] Let J be any $\mathbf{T}$-instance that is a solution for I with respect to $\mathcal{M}$. Take J' to be the $\mathbf{T} \cup \{Eq\}$-instance that extends J with all facts of the form $Eq(a,a)$ with $a \in adom(J)$. It is easy to see that J' is a solution for I with respect to $\mathcal{M}^{\mathrm{Eq}}$, and that, for all UCQs q, we have that $q{\downarrow}(J) = q^{\mathrm{Eq}}{\downarrow}(J')$. Moreover, it is immediate from the construction of J' that there is no pair of distinct constants a, b where $J' \models Eq(a,b)$.

[$\Leftarrow$] Let $\mathbf{f}$ be the collection of function symbols appearing in $\mathcal{M}^{\mathrm{Eq}}$. Let J be a $\mathbf{T} \cup \{\mathrm{Eq}\}$-instance that is a solution for I with respect to $\mathcal{M}^{\mathrm{Eq}}$, such that there is no pair of distinct constants a, b where $J \models \mathrm{Eq}(a, b)$. Note that Eq is an equivalence relation and that each equivalence class contains at most one constant (but possibly many null values). Let $\mathbf{f^0}$ be a witnessing collection of functions, such that $(I, J) \models \mathcal{M}^{\mathrm{Eq}}\ [\mathbf{f} \mapsto \mathbf{f^0}]$. We will construct a $\mathbf{T}$-instance J' and a collection $\mathbf{f^1}$ of functions, as follows:

- For every Eq-equivalence class, choose a single representative member. If an equivalence class contains a constant, we use that constant as the representative member. For every value $u \in adom(J)$, denote by $\pi(u)$ the representative member of the Eq-equivalence class to which u belongs.
- J' contains, for every fact $\mathrm{T}(v_1, \ldots, v_n)$ of J (where $\mathrm{T} \in \mathbf{T}$), the corresponding fact $\mathrm{T}(\pi(v_1), \ldots, \pi(v_n))$
- $\mathbf{f^1}$ contains, for each function f in $\mathbf{f^0}$, the corresponding function f' given by $f'(\mathbf{u}) = \pi(f(\mathbf{u}))$.

By construction, we have that, for any image $q^{\mathrm{Eq}}(\mathbf{a})$ in J of the equality singularization of a conjunctive query $q(\mathbf{x})$, we have an image $q(\mathbf{a})$ in J', and vice versa. This tells us both that J' is a solution for I w.r.t. $\mathcal{M}$ and that for any UCQ q over $\mathbf{T}$, we have $q{\downarrow}(J') = q^{\mathrm{Eq}}{\downarrow}(J)$. Additionally, since each Eq-class is represented by a single member in J', we have that, for each egd in $\sigma \in \Sigma_{\mathrm{t}}$, the fact that J satisfies σ^{Eq} implies that J' satisfies σ. □

The importance of Proposition 6.3 comes from the following observation. Consider any SOTGD+SOTGD schema mapping $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{\mathrm{st}}, \Sigma_{\mathrm{t}})$. Let $\mathbf{f}$ be the collection of all function symbols occurring in SO tgds in $\Sigma_{st} \cup \Sigma_t$. A solution J for a source instance I with respect to $\mathcal{M}$ is said to be a *free* solution if there is a collection of functions $\mathbf{f^0}$ such that $(I, J) \models \Sigma_{st} \cup \Sigma_t\ [\mathbf{f} \mapsto \mathbf{f^0}]$, and such that each function in $\mathbf{f^0}$ is injective and the functions all have mutually disjoint ranges. Equivalently, in a free solution, each value in $adom(J)$ is the denotation of exactly one ground term. If, furthermore, we have that each value in $adom(J)$ is the denotation of a (unique) term of depth k, then we say that J is a *free solution of rank k*.

Proposition 6.4. *Let $\mathcal{M}$ be the equality singularization of a weakly acyclic* EFSOTGD+SOTGD *schema mapping. There is a natural number $k \geq 0$ such that every source instance I has a free universal solution J of rank k.*

Proof. (sketch) Let $\mathbf{f^0}$ be an arbitrary collection of injective and mutually range-disjoint functions. Let J be the result of chasing I with the SO tgds of $\mathcal{M}$ using these functions. A priori, J is potentially infinite. However, we can show that J is always finite, moreover, of finite rank. This is proved by induction: we associate to each position (R, i) (where R is a relation symbol and i an attribute of R) a rank, namely the maximal number of special edges on an incoming path to (R, i) in the dependency graph *times* the maximal depth of a term occurring in the right-hand side of an SO tgd. Then, we can prove by a straightforward induction on k that for all positions (R, i) of rank k, each value in position (R, i) is the denotation of a term of depth at most k. □

It is not hard to see that the same does not hold in the presence of egds.

The above definition of $\mathcal{M}^{Eq}$ applies only to EFSOTGD+(SOTGD, EGD) schema mappings. However, it can be extended to arbitrary SOTGD+(SOTGD, EGD) schema mappings $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{\mathrm{st}}, \Sigma_{\mathrm{t}})$ as follows: from $\mathcal{M}$, we first construct a schema mapping $\mathcal{M}' = (\mathbf{S}, \mathbf{T}', \Sigma_{copy}, \Sigma'_t)$ where $\mathbf{T}' = \mathbf{T} \cup \{R' \mid R \in \mathbf{S}\}$; $\Sigma_{copy} = \{\forall \mathbf{x}(R(\mathbf{x}) \rightarrow R'(\mathbf{x})) \mid R \in \mathbf{S}\}$; and $\Sigma'_t = \Sigma_t \cup \{\sigma' \mid \sigma \in \Sigma_{st}\}$, where σ' is a copy of σ in which every occurrence of a relation $R \in \mathbf{S}$ is replaced by R'. We then define the equality singularization $\mathcal{M}^{Eq}$ to be the equality singulatization of $\mathcal{M}'$. It is easy to see that Proposition 6.3 and Proposition 6.4 then hold true for arbitrary SOTGD+(SOTGD, EGD) schema mappings.

6.3 The Skeleton Rewriting Step

Suppose $\mathcal{M}$ is a weakly acyclic EFSOTGD+EFSOTGD schema mapping, and $\mathcal{M}^{\mathrm{Eq}}$ is the equality singularization of $\mathcal{M}$. Since $\mathcal{M}^{\mathrm{Eq}}$ admits free universal solu-

tions, we can represent the value of every compound term simply by its syntax. This makes it possible to rewrite $\mathcal{M}^{\mathrm{Eq}}$ in such a way that the syntax of compound terms is captured using specialized relations, and constraints with only simple terms.

The *skeleton* of a term is the expression obtained by replacing all constants and variables by $\bullet$, where $\bullet$ is a fixed symbol that is not a function symbol [5]. Thus, for example, the skeleton of $f(g(x,y),z)$ is $f(g(\bullet,\bullet),\bullet)$. The *arity* of a skeleton s, denoted by $arity(s)$, is the number of occurrences of $\bullet$, and the depth of a skeleton is defined in the same way as for terms. If $s, s'_1, \ldots, s'_k$ are skeletons with $arity(s) = k$, then we denote by $s(s'_1, \ldots, s'_k)$ the skeleton of arity $arity(s'_1) + \cdots + arity(s'_k)$ obtained by replacing, for each $i \leq k$, the i-th occurrence of $\bullet$ in s by s'_i.

Definition 6.2. Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{\mathrm{st}}, \Sigma_{\mathrm{t}})$ be a weakly acyclic EFSOTGD+EFSOTGD schema mapping with rank r and whose most deeply nested term has depth d. Let Θ be the set of functions appearing in Σ_{t}. Define the *skeleton rewriting* of $\mathcal{M}$ as the schema mapping $\mathcal{M}^{\mathrm{skel}} = (\mathbf{S}, \mathbf{T}^{\mathrm{skel}}, \Sigma_{\mathrm{st}}^{\mathrm{skel}}, \Sigma_{\mathrm{t}}^{\mathrm{skel}})$, where:

- For every n-ary relation $\mathrm{T} \in \mathbf{T}$, let $\mathbf{T}^{\mathrm{skel}}$ contain all relations of the form $\mathrm{T}_{s_1,\ldots,s_n}$, where $s_1, \ldots s_n$ are skeletons of depth less than or equal to r.
- For every clause $\phi(\mathbf{x}) \to \mathrm{T}(\tau_1, \ldots, \tau_n)$ of a s-t EFSOTGD in Σ_{st}, let $\Sigma_{\mathrm{st}}^{\mathrm{skel}}$ contain the s-t tgd $\phi(\mathbf{x}) \to \mathrm{T}_{s_1,\ldots,s_n}(\bar{\mathbf{x}})$, where $s_1, \ldots, s_n$ are the skeletons for $\tau_1, \ldots, \tau_n$ respectively, and $\bar{\mathbf{x}}$ is the sequence of variables in $\tau_1, \ldots, \tau_n$.
- For every clause $\phi(\mathbf{x}) \to \mathrm{T}(\tau_1, \ldots, \tau_n)$ of a EFSOTGD in Σ_{t} (where $\mathbf{x} = x_1, \ldots, x_m$), let $\Sigma_{\mathrm{t}}^{\mathrm{skel}}$ contain the tgd $\phi_{s_1,\ldots,s_m}(\mathbf{y}_1, \ldots, \mathbf{y}_m) \to \mathrm{T}_{s'_1,\ldots,s'_n}(\bar{\mathbf{y}}_\mathbf{1}, \ldots, \bar{\mathbf{y}}_\mathbf{n})$, where
  - $s_1, \ldots, s_m$ are Θ-skeletons of depth at most $r * d$;
  - each $\mathbf{y}_i$ is a sequence of arity(s_i) fresh variables;
  - $\phi_{s_1,\ldots,s_m}(\mathbf{y}_1, \ldots, \mathbf{y}_m)$ is obtained from ϕ by replacing each atom $\mathrm{R}(x_1, \ldots, x_l)$ by $\mathrm{R}_{s_1,\ldots,s_l}(\mathbf{y}_1, \ldots, \mathbf{y}_l)$;
  - s'_i is a Θ-skeleton of depth at most $r * d$ such that $s'_i = s_k$ (if τ_i is the term x_k) and $s'_i = t_i(s_1, \ldots, s_m)$ (if τ_i is the term $t_i(x_1, \ldots, x_m)$)
  - $\bar{\mathbf{y}}_\mathbf{i} = (\mathbf{y}_{k_1}, \ldots, \mathbf{y}_{k_{arity(s_k)}})$ (if τ_i is the term x_k) and $\bar{\mathbf{y}}_\mathbf{i} = (y_{1_1}, ..., y_{1_{arity(s_1)}}, ..., y_{m_1}, ..., y_{m_{arity(s_m)}})$ (if τ_i is the term $t_i(x_1, \ldots, x_m)$)

In addition, for each conjunctive query $q(\mathbf{x}) = \exists \mathbf{y} \psi(\mathbf{x}, \mathbf{y})$ over $\mathbf{T}$ with $\mathbf{x} = x_1, \ldots, x_n$ and $\mathbf{y} = y_1, \ldots, y_m$, we denote by $q^{\mathrm{skel}}(\mathbf{x})$ the union of conjunctive queries over $\mathbf{T}^{\mathrm{skel}}$ of the form $\exists \mathbf{z}_1 \ldots \mathbf{z}_m \psi_{s_1,\ldots,s_n,s'_1,\ldots,s'_m}(x_1, \ldots, x_n, \mathbf{z}_1, \ldots, \mathbf{z}_m)$, where $s_1 = \ldots = s_n = \bullet$; $s'_1, \ldots, s'_m$ are Θ-skeletons of depth at most r; and each $\mathbf{z}_i$ is a sequence of fresh variables of length $arity(s'_i)$.

For example, consider the EFSOTGD+EFSOTGD schema mapping $\mathcal{M}$ whose constraints are $\exists f \forall x, y \;\; P(x,y) \to Q(f(y), y, y)$ and $\exists g \forall x, y, z \;\; Q(x,y,z) \to Q(x, y, g(x,y))$. Then $\mathcal{M}^{\mathrm{skel}}$ will include the s-t tgd $P(x,y) \to Q_{f(\bullet),\bullet,\bullet}(x,y,y)$, and the target tgds $Q_{\bullet,\bullet,\bullet}(x,y,z) \to Q_{\bullet,\bullet,g(\bullet,\bullet)}(x,y,x,y)$, and $Q_{f(\bullet),\bullet,\bullet}(x,y,z) \to Q_{f(\bullet),\bullet,g(f(\bullet),\bullet)}(x,y,x,y)$. The full skeleton rewriting of our running example schema mapping is given in Figure 12.

Remark. An optimized version of the schema mapping in Figure 12 is shown in Figure 13, based on the simple observation that in Figure 12 none of $\{\mathrm{T}_{\bullet,\bullet}, \mathrm{T}_{f(\bullet,\bullet),\bullet}, \mathrm{T}_{f(\bullet,\bullet),f(\bullet,\bullet)}\}$ appears on the right-hand side of any tgd, and thus the left-hand sides of many tgds cannot be satisfied in any universal solution, and in turn none of $\{\mathrm{Eq}_{\bullet,f(\bullet,\bullet)}, \mathrm{Eq}_{f(\bullet,\bullet),\bullet}\}$ ever appears on the right-hand side of a remaining tgd in which it does not also appear on the left-hand side. We leave development of a principled approach to optimization for future work.

Proposition 6.5. *Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{st}, \Sigma_t)$ be a weakly acyclic* EFSOTGD+EFSOTGD *schema mapping. Let $\mathcal{M}^{\mathrm{skel}}$ be the skeleton rewriting of $\mathcal{M}$. For every UCQ q over $\mathbf{T}$ and for every $\mathbf{S}$-instance I, we have* $\mathrm{certain}(q, I, M) = \mathrm{certain}(q^{\mathrm{skel}}, I, M')$.

Hint. We show that there exists a solution J for I with respect to $\mathcal{M}$ if and only if there exists a solution J' for I with respect to $\mathcal{M}^{\mathrm{skel}}$. Furthermore, J' (respectively J) can be constructed such that for any UCQ q over $\mathbf{T}$, we have $q{\downarrow}(J) = q^{\mathrm{skel}}{\downarrow}(J')$. To construct J from J', we copy every tuple, and use the

skeletons and their arguments to construct the compound terms. To construct J' from J, we copy every tuple, and, using a witnessing collection of functions $\mathbf{f^0}$ such that $(I, J) \models \mathcal{M}\ [\mathbf{f} \mapsto \mathbf{f_0}]$, and such that each null value is the denotation of a unique term of depth at most r. This term gives us both the skeleton and the arguments that belong in J'. □

6.4 Proof of Theorem 6.2

We finally can prove Theorem 6.2 by combining the above results: Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{\mathrm{st}}, \Sigma_{\mathrm{t}})$ be a weakly acyclic SOTGD+(SOTGD, EGD) schema mapping, and let $\hat{\mathcal{M}}$ be the skeleton rewriting of the equality singularization of $\mathcal{M}$, extended with the egd

$$\mathrm{Eq}_{\bullet,\bullet}(x, y) \rightarrow x = y\ .$$

Furthermore, for any UCQ q over $\mathbf{T}$, let $\hat{q}$ be the skeleton rewriting of the equality simulation of q. Then we claim that XR-certain(q, I, M) = XR-certain$(\hat{q}, I, \hat{\mathcal{M}})$.

It suffices to show that, for all source instances I,

1. I has a solution with respect to $\hat{\mathcal{M}}$ if and only if I has a solution with respect to $\mathcal{M}$.
2. if I has a solution with respect to $\mathcal{M}$, then, for all UCQs q over $\mathbf{T}$, certain$(\hat{q}, I, \hat{\mathcal{M}})$ = certain$(q, I, \mathcal{M})$

The first item follows from Proposition 6.3(a) and Proposition 6.5. The second item follows from Proposition 6.3(b) and Proposition 6.5.

6.5 Related Work

Theorem 6.2 allows us to extend the DLP-rewriting technique of Section 5 to GLAV+(WAGLAV, EGD) schema mappings (and, in fact, to weakly acyclic SOTGD+(SOTGD, EGD) schema mappings). The proof is based on a method for eliminating the existentially quantified variables. Others have considered methods for eliminating existential quantifiers from tgds previously, an early example being Duschka and Genesereth's inverse rules algorithm [17] for acyclic LAV rules, which inspired our approach. Krotzsch and Rudolph describe an existentially quantified variable elimination procedure for schema mappings composed of GLAV constraints and relational denial constraints (a subset of denial constraints with no equality or inequality atoms) that are jointly-acyclic (a relaxation of weak acyclicity) in [27]. Their approach is similar to ours in that it creates extra attributes to represent skolem terms in place of existentially quantified variables, but our constraint language includes the additional expressiveness of egds, whose careful handling is a primary concern of our approach. Marnette studied termination of the chase for schema mappings with target constraints in [35], where he introduced the oblivious skolem chase, a modification of the chase procedure in which skolem terms are allowed to appear in instances. A similar procedure was used to prove the correctness of a limited form of skeleton rewriting in [38].

Equality singularization for tgds was introduced in [35], where it was referred to simply as "singularization". In [38], another equality simulation technique was used, based on substitution. In that presentation, the simulation was woven into the skeleton rewriting step.

Theorem 6.2 is related to a result in an unpublished manuscript [36], which can be stated as follows: given any GLAV+WAGLAV schema mapping $\mathcal{M}$ and every conjunctive query q, one can compute a Datalog program that, given any source instance as input, computes the certain answers of q with respect to $\mathcal{M}$. Note that, conceptually, a Datalog program can be viewed as a GAV+GAV schema mapping where the source schema consists of the EDB predicates and the target schema consists of the IDB predicates.

7 Concluding Remarks

In this paper, we introduced the framework of exchange-repairs and explored the XR-certain answers as an alternative non-trivial and meaningful semantics of queries in the context of data exchange. Exchange-repair semantics differ from other proposals for handling inconsistencies in data exchange in that, conceptually, the inconsistencies are repaired at the source rather than the target. This allows the shared origins of target facts to be reflected in the answers to target queries.

This framework brings together data exchange, database repairs, and disjunctive logic programming, thus enhancing the interaction between three different areas of research. Moreover, the results reported here pave the way for using DLP solvers, such as DLV, for query answering under the exchange-repair semantics.

8 Acknowledgements

The research of all authors was partially supported by NSF Grant IIS-1217869. Kolaitis' research was also supported by the project "Handling Uncertainty in Data Intensive Applications" under the program THALES.

References

[1] F. N. Afrati and P. G. Kolaitis. Repair checking in inconsistent databases: algorithms and complexity. In R. Fagin, editor, *ICDT*, volume 361 of *ACM International Conference Proceeding Series*, pages 31–41. ACM, 2009.

[2] M. Alviano, W. Faber, N. Leone, S. Perri, G. Pfeifer, and G. Terracina. The disjunctive datalog system dlv. In *Datalog*, pages 282–301, 2010.

[3] M. Arenas, P. Barceló, L. Libkin, and F. Murlak. *Relational and XML Data Exchange*. Synthesis Lectures on Data Management. Morgan & Claypool Publishers, 2010.

[4] M. Arenas, L. E. Bertossi, and J. Chomicki. Consistent query answers in inconsistent databases. In V. Vianu and C. H. Papadimitriou, editors, *PODS*, pages 68–79. ACM Press, 1999.

[5] M. Arenas, R. Fagin, and A. Nash. Composition with target constraints. *Logical Methods in Computer Science*, 7(3), 2011.

[6] M. Arenas, J. Pérez, J. L. Reutter, and C. Riveros. The language of plain so-tgds: Composition, inversion and structural properties. *J. Comput. Syst. Sci.*, 79(6):763–784, 2013.

[7] L. E. Bertossi. *Database Repairing and Consistent Query Answering*. Synthesis Lectures on Data Management. Morgan & Claypool Publishers, 2011.

[8] L. E. Bertossi, J. Chomicki, A. Cortés-Calabuig, and C. Gutiérrez. Consistent answers from integrated data sources. In T. Andreasen, A. Motro, H. Christiansen, and H. L. Larsen, editors, *FQAS*, volume 2522 of *Lecture Notes in Computer Science*, pages 71–85. Springer, 2002.

[9] M. Bienvenu. On the complexity of consistent query answering in the presence of simple ontologies. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence, July 22-26, 2012, Toronto, Ontario, Canada.*, 2012.

[10] M. Bienvenu, C. Bourgaux, and F. Goasdoué. Querying inconsistent description logic knowledge bases under preferred repair semantics. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, July 27 -31, 2014, Québec City, Québec, Canada.*, pages 996–1002, 2014.

[11] M. Bienvenu and R. Rosati. Tractable approximations of consistent query answering for robust ontology-based data access. In *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013*, 2013.

[12] L. Bravo and L. E. Bertossi. Logic programs for consistently querying data integration systems. In Gottlob and Walsh [23], pages 10–15.

[13] A. Calì, D. Lembo, and R. Rosati. On the decidability and complexity of query answering over inconsistent and incomplete databases. In F. Neven, C. Beeri, and T. Milo, editors, *PODS*, pages 260–271. ACM, 2003.

[14] A. Calì, D. Lembo, and R. Rosati. Query rewriting and answering under constraints in data integration systems. In Gottlob and Walsh [23], pages 16–21.

[15] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, A. Poggi, and R. Rosati. Ontology-based database access. In *Proceedings of the Fifteenth Italian Symposium on Advanced Database Systems, SEBD 2007, 17-20 June 2007, Torre Canne, Fasano, BR, Italy*, pages 324–331, 2007.

[16] J. Chomicki and J. Marcinkowski. Minimal-change integrity maintenance using tuple deletions. *Inf. Comput.*, 197(1-2):90–121, 2005.

[17] O. M. Duschka and M. R. Genesereth. Answering recursive queries using views. In A. O. Mendelzon and Z. M. Özsoyoglu, editors, *PODS*, pages 109–116. ACM Press, 1997.

[18] B. ten Cate, G. Fontaine, and P. G. Kolaitis. On the data complexity of consistent query answering. In A. Deutsch, editor, *ICDT*, pages 22–33. ACM, 2012.

[19] R. Fagin, P. G. Kolaitis, R. J. Miller, and L. Popa. Data exchange: semantics and query answering. *Theor. Comput. Sci.*, 336(1):89–124, 2005.

[20] R. Fagin, P. G. Kolaitis, L. Popa, and W. C. Tan. Composing schema mappings: Second-order dependencies to the rescue. *ACM Trans. Database Syst.*, 30(4):994–1055, 2005.

[21] A. Fuxman and R. J. Miller. First-order query rewriting for inconsistent databases. *J. Comput. Syst. Sci.*, 73(4):610–635, 2007.

[22] M. Gelfond and V. Lifschitz. The stable model semantics for logic programming. In R. A. Kowalski and K. A. Bowen, editors, *ICLP/SLP*, pages 1070–1080. MIT Press, 1988.

[23] G. Gottlob and T. Walsh, editors. *IJCAI-03, Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence, Acapulco, Mexico, August 9-15, 2003*. Morgan Kaufmann, 2003.

[24] G. Grahne and A. Onet. Data correspondence, exchange and repair. In L. Segoufin, editor, *ICDT*, ACM International Conference Proceeding Series, pages 219–230. ACM, 2010.

[25] T. Janhunen and E. Oikarinen. Capturing parallel circumscription with disjunctive logic programs. In J. J. Alferes and J. A. Leite, editors, *JELIA*, volume 3229 of *Lecture Notes in Computer Science*, pages 134–146. Springer, 2004.

[26] P. G. Kolaitis, J. Panttaja, and W. C. Tan. The complexity of data exchange. In S. Vansummeren, editor, *PODS*, pages 30–39. ACM, 2006.

[27] M. Krötzsch and S. Rudolph. Extending decidable existential rules by joining acyclicity and guardedness. In T. Walsh, editor, *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, pages 963–968. IJCAI/AAAI, 2011.

[28] D. Lembo, M. Lenzerini, and R. Rosati. Source inconsistency and incompleteness in data integration. In A. Borgida, D. Calvanese, L. Cholvy, and M. Rousset, editors, *Proceedings of the 9th International Workshop on Knowledge Representation meets Databases (KRDB 2002), Toulouse France, April 21, 2002*, volume 54 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2002.

[29] D. Lembo, M. Lenzerini, R. Rosati, M. Ruzzi, and D. F. Savo. Inconsistency-tolerant semantics for description logics. In *Web Reasoning and Rule Systems - Fourth International Conference, RR 2010, Bressanone/Brixen, Italy, September 22-24, 2010. Proceedings*, pages 103–117, 2010.

[30] D. Lembo and M. Ruzzi. Consistent query answering over description logic ontologies. In M. Marchiori, J. Z. Pan, and C. de Sainte Marie, editors, *Web Reasoning and Rule Systems, First International Conference, RR 2007, Innsbruck , Austria, June 7-8, 2007, Proceedings*, volume 4524 of *Lecture Notes in Computer Science*, pages 194–208. Springer, 2007.

[31] M. Lenzerini. Data integration: A theoretical perspective. In L. Popa, S. Abiteboul, and P. G.

Kolaitis, editors, *PODS*, pages 233–246. ACM, 2002.

[32] N. Leone, G. Pfeifer, W. Faber, T. Eiter, G. Gottlob, S. Perri, and F. Scarcello. The dlv system for knowledge representation and reasoning. *ACM Trans. Comput. Log.*, 7(3):499–562, 2006.

[33] T. Lukasiewicz, M. V. Martinez, A. Pieris, and G. I. Simari. From classical to consistent query answering under existential rules. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA.*, pages 1546–1552, 2015.

[34] M. C. Marileo and L. E. Bertossi. The consistency extractor system: Answer set programs for consistent query answering in databases. *Data Knowl. Eng.*, 69(6):545–572, 2010.

[35] B. Marnette. Generalized schema-mappings: from termination to tractability. In J. Paredaens and J. Su, editors, *PODS*, pages 13–22. ACM, 2009.

[36] B. Marnette. Resolution and datalog rewriting under value invention and equality constraints. *CoRR*, abs/1212.0254, 2012.

[37] R. Rosati. On the complexity of dealing with inconsistency in description logic ontologies. In *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, pages 1057–1062, 2011.

[38] B. ten Cate, R. L. Halpert, and P. G. Kolaitis. Exchange-repairs: Managing inconsistency in data exchange. In R. Kontchakov and M.-L. Mugnier, editors, *RR*, volume 8741 of *Lecture Notes in Computer Science*, pages 140–156. Springer, 2014.

$$
\begin{aligned}
\mathbf{S} &= \{\mathrm{R}\} \\
\mathbf{T} &= \left\{ \mathrm{T}_{\bullet,\bullet}, \mathrm{T}_{\bullet,f(\bullet,\bullet)}, \mathrm{T}_{f(\bullet,\bullet),\bullet}, \mathrm{T}_{f(\bullet,\bullet),f(\bullet,\bullet)}, \mathrm{Eq}_{\bullet,\bullet}, \mathrm{Eq}_{\bullet,f(\bullet,\bullet)}, \mathrm{Eq}_{f(\bullet,\bullet),\bullet}, \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)} \right\} \\
\Sigma_{\mathrm{st}} &= \left\{ \begin{array}{lcl} \mathrm{R}(x,y) & \rightarrow & \mathrm{T}_{\bullet,f(\bullet,\bullet)}(x,x,y) \\ \mathrm{R}(x,y) & \rightarrow & \mathrm{T}_{\bullet,f(\bullet,\bullet)}(y,x,y) \end{array} \right\} \\
\Sigma_{\mathrm{t}}^{\mathrm{skel}} &= \left\{ \begin{array}{l}
\mathrm{T}_{\bullet,\bullet}(x,y) \wedge \mathrm{Eq}_{\bullet,\bullet}(x,x') \wedge \mathrm{T}_{\bullet,\bullet}(x',y') \rightarrow \mathrm{Eq}_{\bullet,\bullet}(y,y') \\
\mathrm{T}_{\bullet,\bullet}(x,y) \wedge \mathrm{Eq}_{\bullet,\bullet}(x,x') \wedge \mathrm{T}_{\bullet,f(\bullet,\bullet)}(x',y_1',y_2') \rightarrow \mathrm{Eq}_{\bullet,f(\bullet,\bullet)}(y,y_1',y_2') \\
\mathrm{T}_{\bullet,\bullet}(x,y) \wedge \mathrm{Eq}_{\bullet,f(\bullet,\bullet)}(x,x_1',x_2') \wedge \mathrm{T}_{f(\bullet,\bullet),\bullet}(x_1',x_2',y') \rightarrow \mathrm{Eq}_{\bullet,\bullet}(y,y') \\
\mathrm{T}_{\bullet,\bullet}(x,y) \wedge \mathrm{Eq}_{\bullet,f(\bullet,\bullet)}(x,x_1',x_2') \wedge \mathrm{T}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_1',x_2',y_1',y_2') \rightarrow \mathrm{Eq}_{\bullet,f(\bullet,\bullet)}(y,y_1',y_2') \\
\mathrm{T}_{\bullet,f(\bullet,\bullet)}(x,y_1,y_2) \wedge \mathrm{Eq}_{\bullet,\bullet}(x,x') \wedge \mathrm{T}_{\bullet,\bullet}(x',y') \rightarrow \mathrm{Eq}_{f(\bullet,\bullet),\bullet}(y_1,y_2,y') \\
\mathrm{T}_{\bullet,f(\bullet,\bullet)}(x,y_1,y_2) \wedge \mathrm{Eq}_{\bullet,\bullet}(x,x') \wedge \mathrm{T}_{\bullet,f(\bullet,\bullet)}(x',y_1',y_2') \rightarrow \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(y_1,y_2,y_1',y_2') \\
\mathrm{T}_{\bullet,f(\bullet,\bullet)}(x,y_1,y_2) \wedge \mathrm{Eq}_{\bullet,f(\bullet,\bullet)}(x,x_1',x_2') \wedge \mathrm{T}_{f(\bullet,\bullet),\bullet}(x_1',x_2',y') \rightarrow \mathrm{Eq}_{f(\bullet,\bullet),\bullet}(y_1,y_2,y') \\
\mathrm{T}_{\bullet,f(\bullet,\bullet)}(x,y_1,y_2) \wedge \mathrm{Eq}_{\bullet,f(\bullet,\bullet)}(x,x_1',x_2') \wedge \mathrm{T}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_1',x_2',y_1',y_2') \rightarrow \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(y_1,y_2,y_1',y_2') \\
\mathrm{T}_{f(\bullet,\bullet),\bullet}(x_1,x_2,y) \wedge \mathrm{Eq}_{f(\bullet,\bullet),\bullet}(x_1,x_2,x') \wedge \mathrm{T}_{\bullet,\bullet}(x',y') \rightarrow \mathrm{Eq}_{\bullet,\bullet}(y,y') \\
\mathrm{T}_{f(\bullet,\bullet),\bullet}(x_1,x_2,y) \wedge \mathrm{Eq}_{f(\bullet,\bullet),\bullet}(x_1,x_2,x') \wedge \mathrm{T}_{\bullet,f(\bullet,\bullet)}(x',y_1',y_2') \rightarrow \mathrm{Eq}_{\bullet,f(\bullet,\bullet)}(y,y_1',y_2') \\
\mathrm{T}_{f(\bullet,\bullet),\bullet}(x_1,x_2,y) \wedge \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_1,x_2,x_1',x_2') \wedge \mathrm{T}_{f(\bullet,\bullet),\bullet}(x_1',x_2',y') \rightarrow \mathrm{Eq}_{\bullet,\bullet}(y,y') \\
\mathrm{T}_{f(\bullet,\bullet),\bullet}(x_1,x_2,y) \wedge \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_1,x_2,x_1',x_2') \wedge \mathrm{T}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_1',x_2',y_1',y_2') \rightarrow \mathrm{Eq}_{\bullet,f(\bullet,\bullet)}(y,y_1',y_2') \\
\mathrm{T}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_1,x_2,y_1,y_2) \wedge \mathrm{Eq}_{f(\bullet,\bullet),\bullet}(x_1,x_2,x') \wedge \mathrm{T}_{\bullet,\bullet}(x',y') \rightarrow \mathrm{Eq}_{f(\bullet,\bullet),\bullet}(y_1,y_2,y') \\
\mathrm{T}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_1,x_2,y_1,y_2) \wedge \mathrm{Eq}_{f(\bullet,\bullet),\bullet}(x_1,x_2,x') \wedge \mathrm{T}_{\bullet,f(\bullet,\bullet)}(x',y_1',y_2') \rightarrow \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(y_1,y_2,y_1',y_2') \\
\mathrm{T}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_1,x_2,y_1,y_2) \wedge \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_1,x_2,x_1',x_2') \wedge \mathrm{T}_{f(\bullet,\bullet),\bullet}(x_1',x_2',y') \rightarrow \mathrm{Eq}_{f(\bullet,\bullet),\bullet}(y_1,y_2,y') \\
\mathrm{T}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_1,x_2,y_1,y_2) \wedge \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_1,x_2,x_1',x_2') \wedge \mathrm{T}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_1',x_2',y_1',y_2') \rightarrow \\
\qquad\qquad \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(y_1,y_2,y_1',y_2') \\
\mathrm{T}_{\bullet,\bullet}(x,y) \rightarrow \mathrm{Eq}_{\bullet,\bullet}(x,x) \\
\mathrm{T}_{\bullet,\bullet}(x,y) \rightarrow \mathrm{Eq}_{\bullet,\bullet}(y,y) \\
\mathrm{T}_{\bullet,f(\bullet,\bullet)}(x,y_1,y_2) \rightarrow \mathrm{Eq}_{\bullet,\bullet}(x,x) \\
\mathrm{T}_{\bullet,f(\bullet,\bullet)}(x,y_1,y_2) \rightarrow \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(y_1,y_2,y_1,y_2) \\
\mathrm{T}_{f(\bullet,\bullet),\bullet}(x_1,x_2,y) \rightarrow \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_1,x_2,x_1,x_2) \\
\mathrm{T}_{f(\bullet,\bullet),\bullet}(x_1,x_2,y) \rightarrow \mathrm{Eq}_{\bullet,\bullet}(y,y) \\
\mathrm{T}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_1,x_2,y_1,y_2) \rightarrow \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_1,x_2,x_1,x_2) \\
\mathrm{T}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_1,x_2,y_1,y_2) \rightarrow \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(y_1,y_2,y_1,y_2) \\
\mathrm{Eq}_{\bullet,\bullet}(x_1,x_2) \rightarrow \mathrm{Eq}_{\bullet,\bullet}(x_2,x_1) \\
\mathrm{Eq}_{\bullet,f(\bullet,\bullet)}(x_1,x_{2_1},x_{2_2}) \rightarrow \mathrm{Eq}_{f(\bullet,\bullet),\bullet}(x_{2_1},x_{2_2},x_1) \\
\mathrm{Eq}_{f(\bullet,\bullet),\bullet}(x_{1_1},x_{1_2},x_2) \rightarrow \mathrm{Eq}_{\bullet,f(\bullet,\bullet)}(x_2,x_{1_1},x_{1_2}) \\
\mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_{1_1},x_{1_2},x_{2_1},x_{2_2}) \rightarrow \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_{2_1},x_{2_2},x_{1_1},x_{1_2}) \\
\mathrm{Eq}_{\bullet,\bullet}(x_1,x_2) \wedge \mathrm{Eq}_{\bullet,\bullet}(x_2,x_3) \rightarrow \mathrm{Eq}_{\bullet,\bullet}(x_1,x_3) \\
\mathrm{Eq}_{\bullet,\bullet}(x_1,x_2) \wedge \mathrm{Eq}_{\bullet,f(\bullet,\bullet)}(x_2,x_{3_1},x_{3_2}) \rightarrow \mathrm{Eq}_{\bullet,f(\bullet,\bullet)}(x_1,x_{3_1},x_{3_2}) \\
\mathrm{Eq}_{\bullet,f(\bullet,\bullet)}(x_1,x_{2_1},x_{2_2}) \wedge \mathrm{Eq}_{f(\bullet,\bullet),\bullet}(x_{2_1},x_{2_2},x_3) \rightarrow \mathrm{Eq}_{\bullet,\bullet}(x_1,x_3) \\
\mathrm{Eq}_{\bullet,f(\bullet,\bullet)}(x_1,x_{2_1},x_{2_2}) \wedge \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_{2_1},x_{2_2},x_{3_1},x_{3_2}) \rightarrow \mathrm{Eq}_{\bullet,f(\bullet,\bullet)}(x_1,x_{3_1},x_{3_2}) \\
\mathrm{Eq}_{f(\bullet,\bullet),\bullet}(x_{1_1},x_{1_2},x_2) \wedge \mathrm{Eq}_{\bullet,\bullet}(x_2,x_3) \rightarrow \mathrm{Eq}_{f(\bullet,\bullet),\bullet}(x_{1_1},x_{1_2},x_3) \\
\mathrm{Eq}_{f(\bullet,\bullet),\bullet}(x_{1_1},x_{1_2},x_2) \wedge \mathrm{Eq}_{\bullet,f(\bullet,\bullet)}(x_2,x_{3_1},x_{3_2}) \rightarrow \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_{1_1},x_{1_2},x_{3_1},x_{3_2}) \\
\mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_{1_1},x_{1_2},x_{2_1},x_{2_2}) \wedge \mathrm{Eq}_{f(\bullet,\bullet),\bullet}(x_{2_1},x_{2_2},x_3) \rightarrow \mathrm{Eq}_{f(\bullet,\bullet),\bullet}(x_{1_1},x_{1_2},x_3) \\
\mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_{1_1},x_{1_2},x_{2_1},x_{2_2}) \wedge \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_{2_1},x_{2_2},x_{3_1},x_{3_2}) \rightarrow \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_{1_1},x_{1_2},x_{3_1},x_{3_2})
\end{array} \right\} \\
\mathrm{reachable}^{\mathrm{skel}}(x,y) &:= \exists c,c'(\mathrm{T}_{\bullet,\bullet}(x,c) \wedge \mathrm{Eq}_{\bullet,\bullet}(c,c') \wedge \mathrm{T}_{\bullet,\bullet}(y,c')) \\
&\cup \exists c,c_1',c_2'(\mathrm{T}_{\bullet,\bullet}(x,c) \wedge \mathrm{Eq}_{\bullet,f(\bullet,\bullet)}(c,c_1',c_2') \wedge \mathrm{T}_{\bullet,f(\bullet,\bullet)}(y,c_1',c_2')) \\
&\cup \exists c_1,c_2,c'(\mathrm{T}_{\bullet,f(\bullet,\bullet)}(x,c_1,c_2) \wedge \mathrm{Eq}_{f(\bullet,\bullet),\bullet}(c_1,c_2,c') \wedge \mathrm{T}_{\bullet,\bullet}(y,c')) \\
&\cup \exists c_1,c_2,c_1',c_2'(\mathrm{T}_{\bullet,f(\bullet,\bullet)}(x,c_1,c_2) \wedge \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(c_1,c_2,c_1',c_2') \wedge \mathrm{T}_{\bullet,f(\bullet,\bullet)}(y,c_1',c_2'))
\end{aligned}
$$

Figure 12: Undirected reachability, skolemized, equality singularized, and skeleton rewritten.

$$
\begin{array}{ll}
\mathbf{S} & = \{\mathrm{R}\} \\
\mathbf{T} & = \left\{ \mathrm{T}_{\bullet,f(\bullet,\bullet)}, \mathrm{Eq}_{\bullet,\bullet}, \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)} \right\} \\
\Sigma_{\mathrm{st}} & = \left\{ \begin{array}{lll} \mathrm{R}(x,y) & \rightarrow & \mathrm{T}_{\bullet,f(\bullet,\bullet)}(x,x,y) \\ \mathrm{R}(x,y) & \rightarrow & \mathrm{T}_{\bullet,f(\bullet,\bullet)}(y,x,y) \end{array} \right\} \\
\Sigma_{\mathrm{t}}^{\mathrm{skel}} & = \left\{ \begin{array}{l} \mathrm{T}_{\bullet,f(\bullet,\bullet)}(x,y_1,y_2) \wedge \mathrm{Eq}_{\bullet,\bullet}(x,x') \wedge \mathrm{T}_{\bullet,f(\bullet,\bullet)}(x',y_1',y_2') \rightarrow \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(y_1,y_2,y_1',y_2') \\ \mathrm{T}_{\bullet,f(\bullet,\bullet)}(x,y_1,y_2) \rightarrow \mathrm{Eq}_{\bullet,\bullet}(x,x) \\ \mathrm{T}_{\bullet,f(\bullet,\bullet)}(x,y_1,y_2) \rightarrow \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(y_1,y_2,y_1,y_2) \\ \mathrm{Eq}_{\bullet,\bullet}(x_1,x_2) \rightarrow \mathrm{Eq}_{\bullet,\bullet}(x_2,x_1) \\ \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_{1_1},x_{1_2},x_{2_1},x_{2_2}) \rightarrow \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_{2_1},x_{2_2},x_{1_1},x_{1_2}) \\ \mathrm{Eq}_{\bullet,\bullet}(x_1,x_2) \wedge \mathrm{Eq}_{\bullet,\bullet}(x_2,x_3) \rightarrow \mathrm{Eq}_{\bullet,\bullet}(x_1,x_3) \\ \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_{1_1},x_{1_2},x_{2_1},x_{2_2}) \wedge \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_{2_1},x_{2_2},x_{3_1},x_{3_2}) \rightarrow \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(x_{1_1},x_{1_2},x_{3_1},x_{3_2}) \end{array} \right\} \\
\mathrm{reachable}^{\mathrm{skel}}(x,y) & := \quad \exists c_1,c_2,c_1',c_2'(\mathrm{T}_{\bullet,f(\bullet,\bullet)}(x,c_1,c_2) \wedge \mathrm{Eq}_{f(\bullet,\bullet),f(\bullet,\bullet)}(c_1,c_2,c_1',c_2') \wedge \mathrm{T}_{\bullet,f(\bullet,\bullet)}(y,c_1',c_2'))
\end{array}
$$

Figure 13: Example schema mapping from Figure 9, skolemized, equality singularized, skeleton rewritten, and optimized.